AD-753 656

ADVANCED RESEARCH PROJECTS AGENCY (ARPA)
SEISMIC COUPLING CONFERENCE HELD AT INSTI-
TUTE FOR DEFENSE ANALYSES, ARLINGTON,
VIRGINIA, JUNE 8-9, 1970

Battelle Columbus Laboratories

Prepared for:

Advanced Research Projects Agency

1972

ARPA-TIO -71-13-1

ARPA SEISMIC COUPLING CONFERENCE

Held at Institute for Defense Analyses
Arlington, Virginia

June 8-9, 1970

D D C
RECEIVED
AUG 21 19..
D..

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

ADVANCED RESEARCH PROJECTS AGENCY

1400 WILSON BLVD
ARLINGTON, VA. 22209

DOCUMENT CONTROL DATA - R&D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) BATTELLE Columbus Laboratories 505 King Avenue Columbus, Ohio 43201	2a. REPORT SECURITY CLASSIFICATION Unclassified
	2b GROUP

3. REPORT TITLE

ARPA SEISMIC COUPLING CONFERENCE.
Held at Institute for Defense Analysis, Arlington, Virginia, June 8-9, 1970.

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

5. AUTHOR(S) (Last name, first name, initial)

6. REPORT DATE Published 1972	7a. TOTAL NO. OF PAGES ~~269~~ 266	7b. NO. OF REFS 106
8a. CONTRACT OR GRANT NO. DAHC-15-70-C-0259, Mod. P 00003	9a. ORIGINATOR'S REPORT NUMBER(S) ARPA-TIO-71-13-1	
b. PROJECT NO.		
c. ARPA Order No. 1594	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)	
d.		

10 AVAILABILITY/LIMITATION NOTICES

Approved for public release; distribution unlimited

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency

13. ABSTRACT

This conference and a subsequent one in August 1970 (reported in ARPA-TIO-71-13-2) were held to foster communication among the diverse disciplines required to predict the shock effects from nuclear explosions out to teleseismic distances. These disciplines involve the use of rock mechanics, geology, nuclear physics, computer hardware and codes, seismology, and field instrumentation. Results from the conferences included (a) improvement in the communication links between the engineers and scientists engaged in research relevant to the seismic coupling problems, and (b) identification of open circuits at some points along the communication lines. This report presents the June 1970 conference proceedings and a summary paper on the results of both conferences.

DD FORM 1473
1 JAN 64

14.		LINK A		LINK B		LINK C	
KEY WORDS		ROLE	WT	ROLE	WT	ROLE	WT
Seismic detection							
Nuclear explosion detection							

1h

ARPA SEISMIC COUPLING CONFERENCE

Held at
Institute for Defense Analysis
Arlington, Virginia

June 8-9, 1970

Proceedings prepared by

BATTELLE
Columbus Laboratories
505 King Avenue
Columbus, Ohio 43201

TABLE OF CONTENTS

ii

FIGURES (Continued)

FIGURES (Continued)

FIGURES (Continued)

WELCOME AND INTRODUCTORY REMARKS

S. J. Lukasik
ARPA

Good morning. I would like to welcome you. We are very happy to have all of you here. I am, of course, pleased to see the amount of interest in the technical community on our subject.

Let me outline the extent of the ARPA interest in this area so that we can better understand each other's motivations.

This meeting is the second of its kind. The first meeting was held a little over two years ago, in January '68. At that time we believed there was a lack of communication between the people who did multidimensional hydrodynamics calculations, the people who understood the properties of rocks, particularly real materials as they occur in bulk, and the people who are involved in various sorts of problems related to solid earth geophysics having a military interest. I can mention two widely different kinds of solid earth geophysicists as examples. One is those who are interested in what went on at the Rocky Mountain Arsenal and the apparent generation of earthquakes as a result of fluid injection. Another example is the community involved in calculating underground explosions.

The ARPA interest in this latter field has been related to the work in our Nuclear Monitoring Research Office on seismic detection and identification for the purpose of supporting negotiations on underground test ban treaties.

There are several reasons why we are interested in such phenomena. In particular we have been concerned with what someone can do to evade a treaty. One must think about the problems that a treaty should address: what you can do to detect and identify possible violations and what the other side can do to circumvent your doing this.

This consideration has motivated us since about 1965. We have pursued this question largely in an empirical way; the reason being that the arithmetic is complicated and many of the important physical properties of the systems involved are not known. The usual way to proceed in such a case is to take an empirical approach, which means firing underground shots or doing some sort of scaled experiments. Because we do have a fairly broad interest in this, let me mention several particular things that have concerned us.

First, there is the coupling of underground explosions in porous media. It is easy to talk about granite, but quite a different thing to discuss porous media--either alluvium or porous mixtures of rocks. All sorts of tuffs occur in a wide variety of sizes and shapes and have different properties. The problem of calculating how much

1

seismic signal comes from a given yield in such surroundings is
fairly complicated.

As another example, utilizing the so-called cavity decoupling
concept one can fire nuclear devices in holes. But when you look into
this possibility you find that the holes are large, expensive, and
perhaps not credible. On the other hand, it turns out you can overdrive
and fire larger yields than one would expect for full decoupling without
paying too high a price in seismic signal. The details of what happens
in such a case depend critically on the behavior of rock stressed beyond
its elastic limit which means phase changes, cracking, and so on.

The whole business of the construction of cavities in various
media requires an understanding of the properties of rock and involves
fairly complicated calculations. We have had occasion to look into the
question of the stability of such cavities because one of the questions
immediately asked is "How big a cavity can you make?" Depending on
what you do thereafter, the answer essentially sets an upper limit on
the yield one can possibly use and still get relatively small signals

Finally, there is another scheme we have been interested in--
the possibility of hiding a clandestine shot in an earthquake. At
first glance this procedure doesn't seem enormously useful because
earthquakes of the right size and location are not that frequent, and
it doesn't seem like a very practical way to run a weapon test program.
On the other hand, if one has control of the earthquake one is hiding
in, that fact changes the likelihood by several orders of magnitude.
Thus we are interested in the question of how one creates earthquakes.
From the Rocky Mountain Arsenal experience it would appear that it is
possible to generate earthquakes.

Thus in all of these problem areas there is a need for under-
standing the properties of the material under explosive loading con-
ditions as well as on the slower time scale involved in hydrofracturing
or some other technique for relieving stresses and causing changes in
natural seismicity.

So you see why we are rather interested in this whole subject
area. But what is the point of getting everyone together? Just what
are we pushing in particular? At the risk of displaying a personal
prejudice, it seems to me that in the long run the class of questions
that are involved is very large. The cost of approaching all of these
problems empirically is very high. The number of questions and the
degree to which we want answers probably exceeds our pocketbook, and
at the present rate it appears that they will become even more so in
the future. Therefore, it would seem that we should proceed in a
different manner. It is my belief, and I am sure it is shared by others
here, that one can calculate the answers to a large number of the
questions. I won't be so rash as to say one can calculate sufficient
answers to all the questions involved, and I think that is one of the

2

points of a meeting such as this, to help clarify what the balance of computation, laboratory work, and field experiments should be.

In any event, if one is going to do an experiment, one can accomplish much by computation in searching through the many parameters that are involved to select the best experiment. In so far as one is given an unlimited computer capacity one can do almost anything. The key point is computer capacity.

Several things have happened in the last few years since we started thinking about this aspect. There is, of course, an impressive amount of computer capacity in the world already, and you have all seen some of the results. On the other hand, I think that when one looks carefully into this, one concludes that the details of the questions we can ask and the number of such questions probably exceed the computer capacity available today, at least in terms of any realistic computing budget. When the computing job looks as though it will cost one or two million dollars, you begin thinking about doing the nuclear shot which costs about the same amount and doesn't have the theoretical ambiguities associated with it.

We have been working with that problem in our Information Processing Techniques Office. For several years they have supported the construction of a high-speed parallel processor called ILLIAC IV that Burroughs is building under subcontract to the University of Illinois. This is, of course, not the only supercomputer which is under construction, and I am sure that other computer companies have products in design. But it is always easier to talk about something that you know about, is likely to come into existence, and we already own anyway. It is a very fast machine. We expect the ILLIAC IV to be faster by a significant amount than any existing or planned machine and to be on line sometime in 1971. A substantial fraction of this machine has been allocated to problems of the kind I mentioned earlier.

Therefore, I think that comments like "That is all very well, but we don't have enough computing capacity" are not adequate reason to preclude the computational approach to the various problems that are the subject of today's discussions.

Data to be input to the machine is another question. It is somewhat indefinite. I hope that one of the primary results of this and presumably subsequent meetings would be to settle some of these difficulties. We all have used computers enough to know one of the first rules of computation is "Garbage in, garbage out." It is very easy to generate garbage; computers are very good at that. Therefore, it is extremely important that we understand the physics of our processes when we write the programs and that we understand the material properties that we input to the machine.

This is now the key point in this whole program. I think we know _what_ we want to do. I think we know _how_ to write hydrodynamic

3

codes. We have or soon will have adequate computer capacity, at least in order to go on to the next step. But I am not at all certain that we _understand_ the properties of the materials, either in general or as far as the specific numbers that should be used to represent a given material under pertinent physical conditions.

To broaden the subject somewhat, I would like to point out one other thing that has happened since the first meeting several years ago. ARPA, in our Nuclear Monitoring Research Office, has started a new program called the Military Geophysics Program. While it is related to, and stimulated by, much of the physics that has been involved in the treaty evasion--Project-VELA-kinds of interests, it is a new program and it is intended to explore the military and defense consequences of a broader range of problems in solid earth geophysics. One of them I have mentioned already, the generation of earthquakes, at Rocky Mountain Arsenal. This is not a direct concern of Project VELA, although it does stimulate us to think about what would happen if one could generate earthquakes on demand. Another problem of concern is seismicity related to large underground explosions.

Finally, there is a part of the program I would like to just mention although it is not a part of the program in the Defense Department at this time. That is the possibility of earthquake control. It is likely that if one understands enough about the properties of materials and one does just the right things, one can relieve crustal stress in a way that is less damaging than what happens when there is a sudden release of stress energy in the form of a large earthquake.

We have also looked into the problems of underground excavation and the appropriate properties of rocks. I mention for completeness that we are interested in instrumentation, the problems of measuring properties particularly of rocks at great depths that one needs use in the computer codes. We are interested in a number of similar things, problems that are fundamentally amenable to calculation but ones where we must have a good understanding of the properties and materials as well as having a rather substantial computational capability to attack realistic problems in a realistic way.

Well, I think this introduction is enough. My only purpose has been to indicate to you some idea of the extent and breadth of our interest. Let me say again that I am very pleased to have you all here. I hope the discussions will be fruitful and that we will come to some general agreement as to what can be done and what can't be done and what the next steps are.

ROCK PROPERTIES

John Handin
Texas A&M University

I would like to begin by philosophizing a little about rock-properties testing, then briefly to review what we know about the constitutive relations of rocks under conditions relevant to the seismic-coupling problem, and finally to suggest where more work is needed.

The measurement of rock properties has a kind of uncertainty principle of its own. How do we select and collect samples that are both undisturbed and representative?

To obtain a laboratory-size specimen we must forcefully remove an element of the material from its natural environment. In contrast to soils, rocks can usually be sampled without seriously disturbing their structures if due care is taken. The question is, do we properly simulate the natural environment when we return the rock to the laboratory and test it? I think the answer is affirmative, provided that the special boundary conditions imposed by any particular test are accounted for.

The choice of samples is much more difficult. What is a so-called "representative" sample of the real rock mass? Rock is structurally and compositionally complex. It is never ideally (and rarely even statistically) homogeneous and isotropic, especially in domains of a few cubic meters or more. The shallow crust is characterized by rapid lithologic variations both vertically and laterally, and because of very nearly ubiquitous layering and fracturing (or jointing), it cannot be validly treated as a mechanical continuum. Are laboratory tests on a few small samples from a few points in the real heterogeneous rock mass worth doing at all? I say yes, stipulating that the testing is done for the right reasons, because the only alternative, in situ testing, is now very expensive and time consuming. Furthermore, current field measurements are usually, if not always, subject to uncertain interpretation, and their extrapolations to neighboring regions of the rock mass are nearly as risky as are those of the laboratory measurements themselves.

What are the right reasons for laboratory testing? The all-too-common, blindly empirical collection of raw data is surely not one of them. Rather the principal purpose should be to provide realistic working models of rock masses--the mathematical models of the media that are plugged into the code calculations to predict ground motions. We have long had enough experience to place reasonable upper and lower bounds on the response to dynamic loading, given even a limited sampling. If, for example, the Piledriver rock had been treated not as a strong granite but as the blocky, water-saturated mass geologists knew it to

5

be, the recorded ground motions need have surprised no one. Although a general theory of rock failure is unlikely to emerge soon, laboratory work can provide empirical models which can in turn help guide the vastly more expensive field tests that will be needed for the ultimate check on our concepts of seismic coupling.

The constitutive relations of rocks depend on a baffling number of intrinsic and extrinsic variables--at least on composition and fabric, the state of effective stress, temperature, strain rate, and strain too because of work hardening or softening. The purpose of laboratory testing is not only to delineate the modes of failure and measure the mechanical properties as functions of these variables, but to sort out the parameters and assess their relative importance. Even if a completely realistic failure criterion could be devised, it would not be mathematically tractable. Our working models will necessarily always be idealizations of the real rock mass. Laboratory experiments that simulate the natural environment as realistically as possible and that reckon favorably with the facts of nature can help us select the successful models that most nearly approximate the in situ deformations but are also practicable. The choice should be made in the light of the geological description and of the parametric studies of code calculations.

It has been suggested that as long as the purely mathematical theories provide solutions fitting the recorded ground motions, then the physics of the deformation is only of academic interest. Indeed, using the Drucker-Prager yield condition and associated flow rules, we have rather successfully predicted the deformations of loose, granular, triboplastic media like alluvium without really knowing why. However, I reject this approach to problems in rock mechanics. Ignoring the geology and ignorance of the physics are exactly the reasons for our severely limited ability to predict ground motions in untested rock media.

I would also like to ask for standardization, without which interlaboratory comparisons of data will always be suspect. Standardization involves both the testing procedure and the material itself. The ASTM has published (or soon will issue) standards for static, uniaxial compression and tensile tests, for dynamic measurements of elastic properties, and for static, triaxial-compression testing under confining pressures to about 1 kb. One might prefer somewhat different procedures. Nevertheless I recommend the universal adoption of these ASTM standards. At least we can then be sure all data are comparable. Experiments at high confining pressure are done pretty much according to individual taste. Interlaboratory correlations are not valid unless the peculiar constraints of each apparatus are duly accounted for.

What about the material? Many interlaboratory studies have been and doubtless will be made of what is supposed to be the same rock. Considering heterogeneity of rock, we can judge that samples are identical only by the careful petrographic examination that is too often neglected.

6

At the last ARPA seismic-coupling meeting in Menlo Park, about 2-1/2 years ago, the discussion centered on the best choice of failure criteria for intact rock. Although the physics of the fracture of brittle rock and the so-called "plastic" flow of ductile rock were not well known (and still aren't), there seemed to be general agreement that the Coulomb-Mohr fracture criterion and the Mises yield condition were at least adequate mathematical formalisms. Several of the participants emphasized, however, that the real rock mass is not a homogeneous, isotropic continuum and that to predict the deformation one must know not only the properties of the intact rocks but also the influence of defects such as layering and jointing of which few rocks are free.

Although we still need more information on dynamic strength and on unloading--that is, the complete stress-strain curve, both dynamic and static--our knowledge of the mechanical properties of intact, homogeneous rock is pretty good. This information is well known to some of us, but the rock-properties people do not seem to have effectively communicated their ideas to the seismic-coupling community, so let me review these data briefly by examining some typical stress-strain curves.

At constant temperature and strain rate (Figure 1A), effective confining pressure, increasing upward in curves A through D, enhances ultimate strength or peak stress difference (curve B), raises the yield stress (curve C), and favors work- or strain-hardening (curve D). Ductility is also enhanced as the rock tends to pass from the brittle state (curve A) through transitional states (curves B and C) to the fully ductile state of uniform flow (curve D).

At constant effective confining pressure (Figure 1B), increasing the temperature at constant strain rate, or what is equivalent, decreasing the strain rate at constant temperature both tend to reduce ultimate strength and yield stress, to increase ductility (curves E through H), and to eliminate strain hardening.

According to Terzaghi (1943)*, the effective stress is the difference between the total stress and the hydrostatic pore pressure; it is the stress in the solid framework of the rock.

Let's consider the influence of pore pressure on the inelastic deformation of a porous, permeable sandstone (Figure 2). The lower series of stress-strain curves shows two important effects of increasing the pore pressure (or decreasing the effective pressure) from 0 to 2 kb at a constant total confining pressure of 2 kb. Clearly ultimate strength and ductility are both reduced. The well defined peaks of the lower three curves signal the onset of faulting.

_____

* References are listed alphabetically on pages 30-31.

7
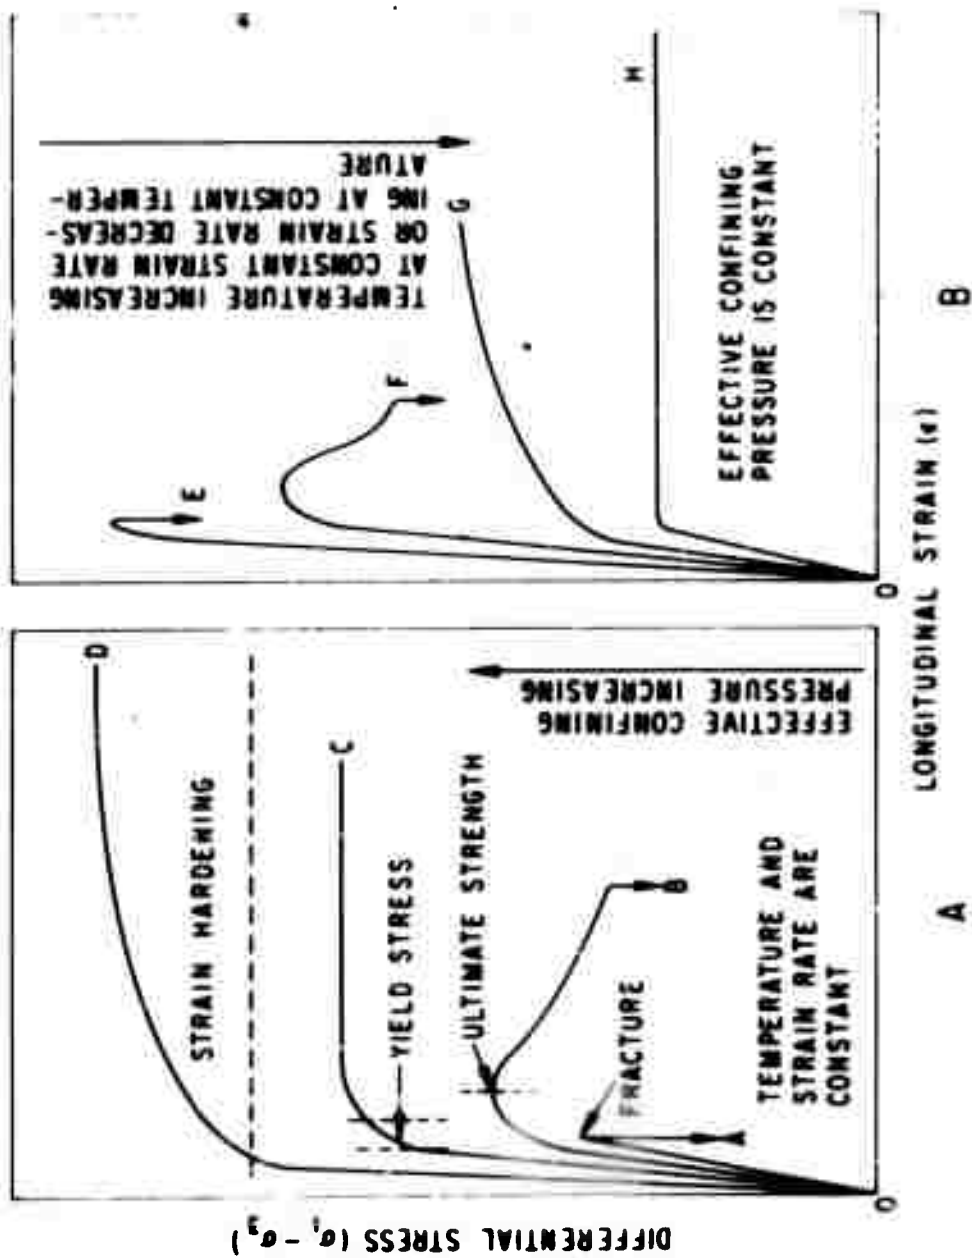
Figure 1. Typical Stress-Strain Curves for Rocks Showing the Effects of Temperature, Strain Rate, and Effective Confining Pressure.

8
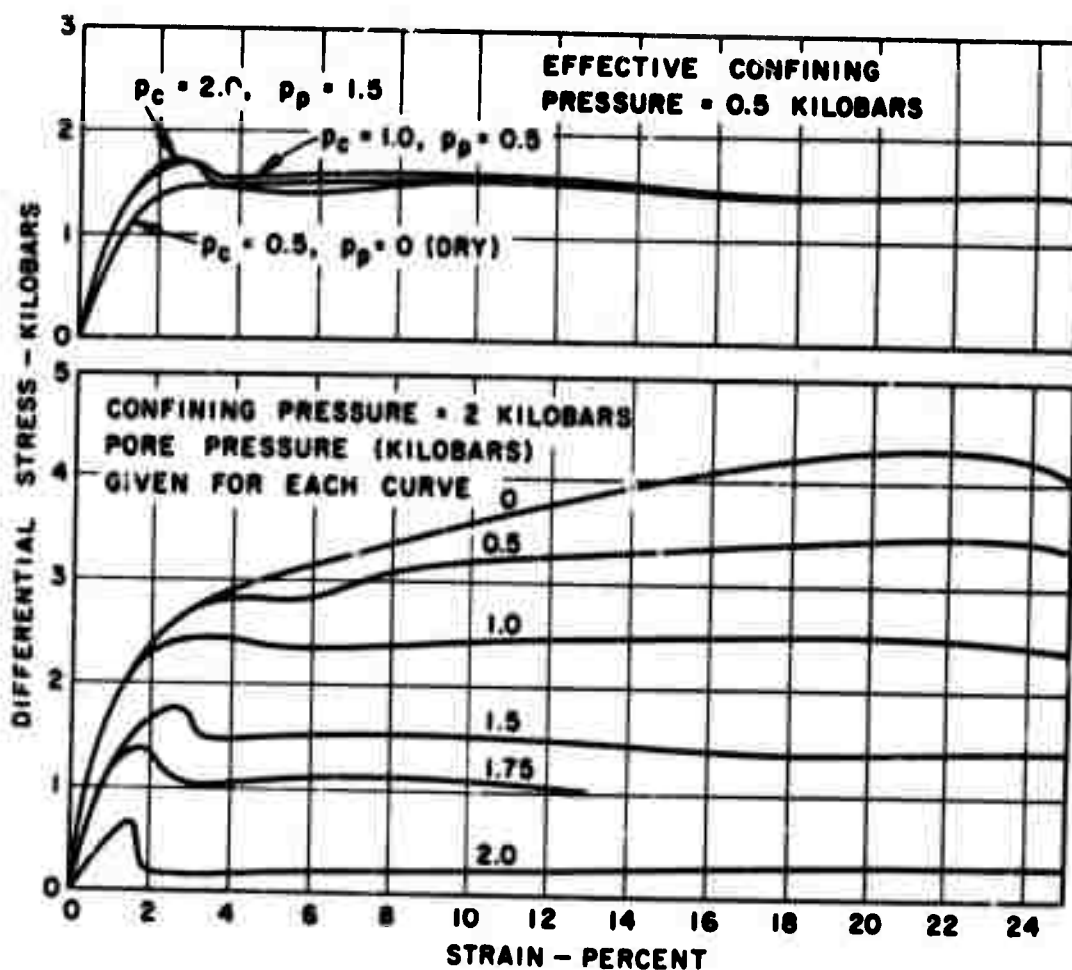
Figure 2. Stress-Strain Curves for Berea Sandstone Compressed at Different Pore-Water Pressures. p_C = confining pressure; p_p = pore pressure; $p_C - p_p$ = effective confining pressure. (After Handin et al., 1963, Figure 4).

9

That these properties are functions only of effective stresses and not of the absolute values of either total or pore pressures is evident from comparisons of the three essentially identical curves in upper part of Figure 2. This rock behaves similarly in every test in which the effective pressure is the same, namely, 0.5 kb.

Another way of demonstrating this principle is by drawing the Mohr envelopes. One-half of the stress circles representing the extreme principal stresses at the failure of the dry rock are constructed in the lower part of Figure 3.
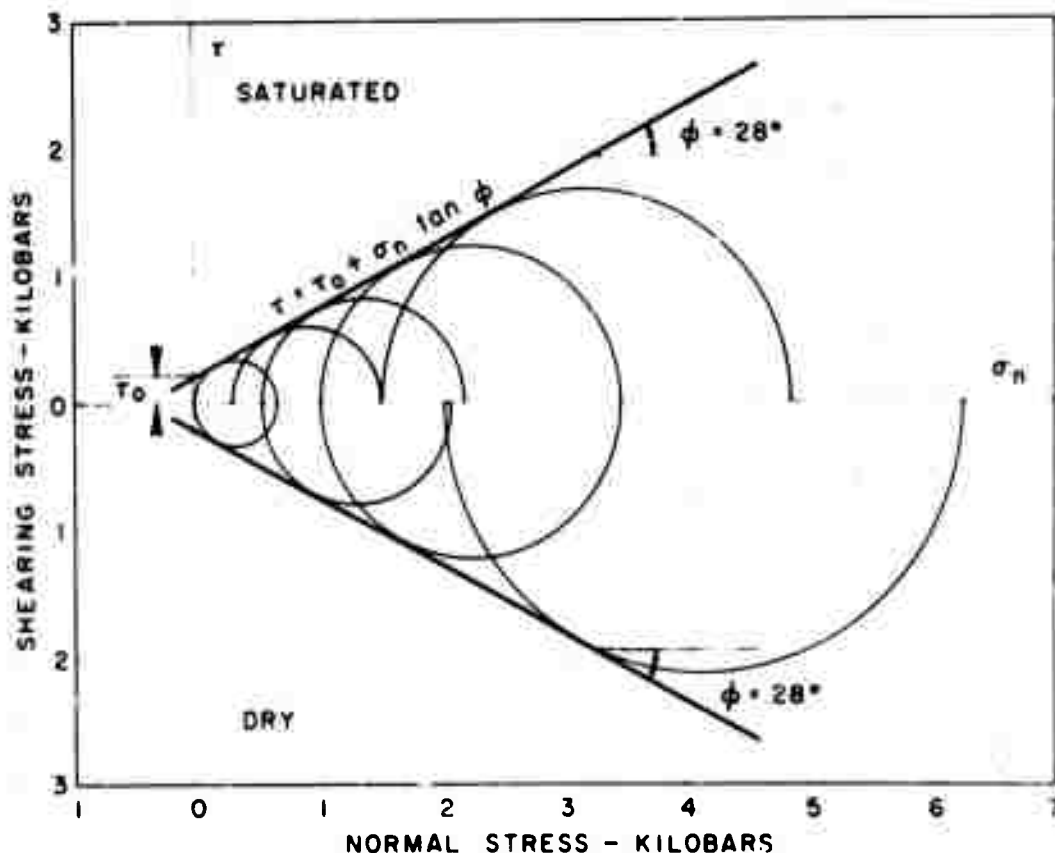
Figure 3. Mohr Envelopes (Identical) for the Ultimate Compressive Strength of Dry and Water-Saturated Berea Sandstone. (After Handin et al., 1963, Figure 7).

The linear envelope tangent to these circles represents the Coulomb criterion where the cohesive resistance at zero normal stress is 200 bars, and the angle of internal friction is a little less than 30 deg. The envelope tangent to the principal effective stress circles at the ultimate compressive strength of the water-saturated rock, plotted in the upper half of the figure, is virtually identical. The center of a stress circle merely moves toward zero along the normal-stress axis by the amount of the pore pressure. Note that if the cohesive strength were zero, that is, if the rock contained a pre-existing fracture, the

10

envelopes would pass through the origin of the τ-σ diagram, and when the total and pore pressures were equal, the rock would possess vanishing shear strength.

Brace (1968) has shown that the effective stress concept is even applicable to rocks like granite with extremely low porosities and permeabilities, provided that the strain rate is low enough for equilibration of pore pressure. If the time of compressive strain is shorter than the equilibration time, there are two possibilities. In rocks with porosities of a few percent or more, the void volume may decrease. The pore pressure will then rise, and the strength will decrease (Handin et al., 1963).

On the other hand, compact crystalline rocks like granite and basalt may be dilatant even under confining pressures of several kilobars. Idealized stress-strain curves reflect four regions of deformation (Figure 4). In region 1 where the stress level is low, the axial and volumetric stress-strain curves are both nonlinear, owing presumably to the closing of microcracks. This is also the region where dynamic and static measurements of elastic "constants" fail to agree. In region 2 pre-existing cracks have closed, and both curves are linear-- that is, the deformation is perfectly elastic. Volume is decreasing. In high-stress region 3 the axial strain still increases, though inelastically. However, the volume strain reverses sign. It increases as the rock becomes dilatant, owing presumably to the opening of microcracks parallel to the maximum principal compressive stress. Here pore pressure will fall and strength will increase. Brace (1968) appropriately calls this phenomenon dilatancy hardening. Both hardening and softening should be looked for in dynamic tests on saturated rocks.

Dilatancy is premonitory to macroscopic shear fracturing after which the deformation in region 4 is due to slip on the new fracture surface.

Scholz (1968) has confirmed these suggested mechanisms by listening to the sounds emitted by deforming specimens (Figure 5). Under the high confining pressure of 4 kb most pre-existing cracks have already closed so that few microseismic events occur at low-stress levels. Beginning at a differential stress of about 8 kb, half of the macroscopic breaking strength of this Westerly granite specimen, the number of events increases rapidly and goes off scale just prior to rupture.

Swanson (1969) has also observed dilatancy, and he has shown how to incorporate this effect into a constitutive model. Curves of shear stress versus shear strain and mean pressure versus volume strain in cyclically loaded Westerly granite reflect significant hysteresis (Figure 6). Porous rocks can be hysteretic even under purely hydrostatic stress if their frameworks break down.

11
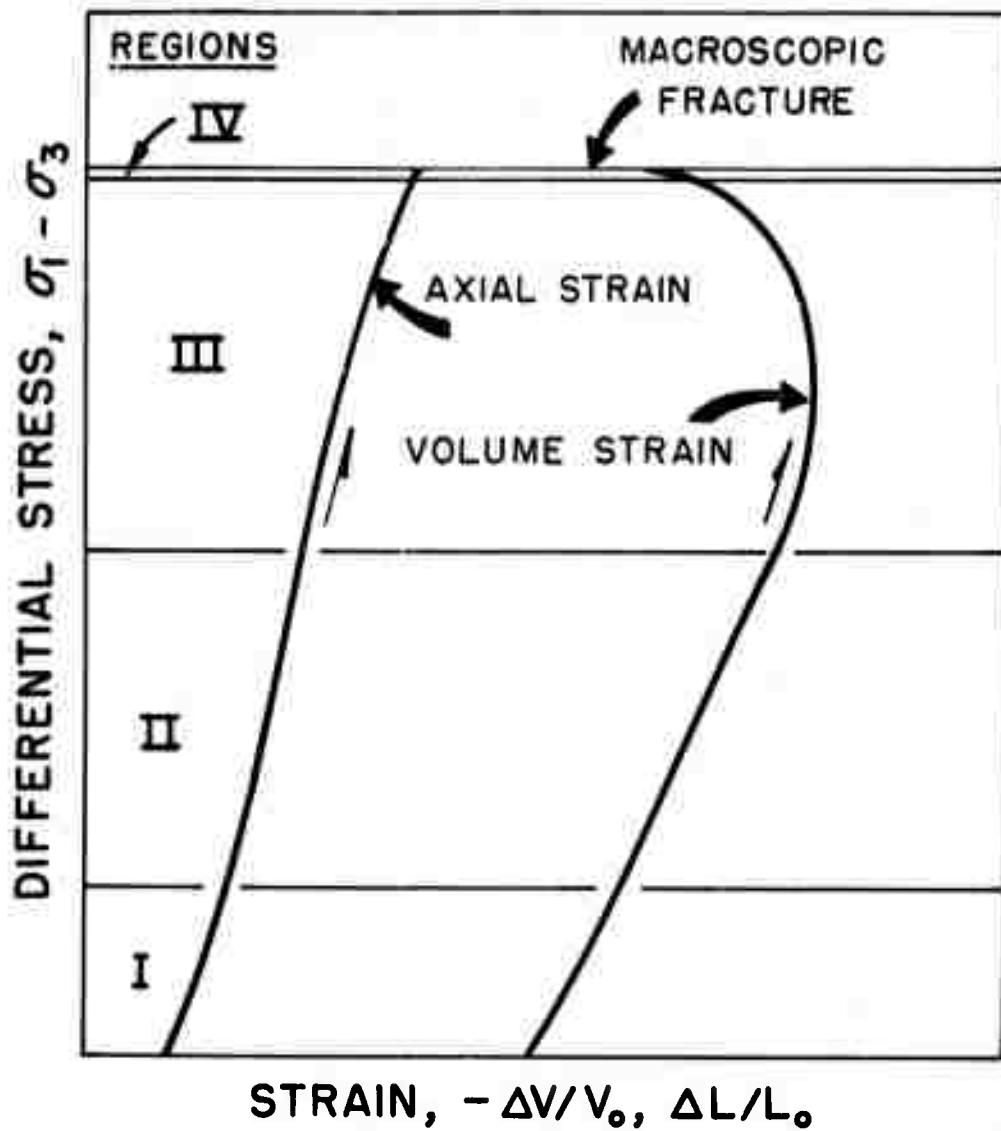
Figure 4. Idealized Triaxial Compression Stress-Strain Curves for Compact Crystalline Rock. (After Brace et al., 1966, Figure 5).
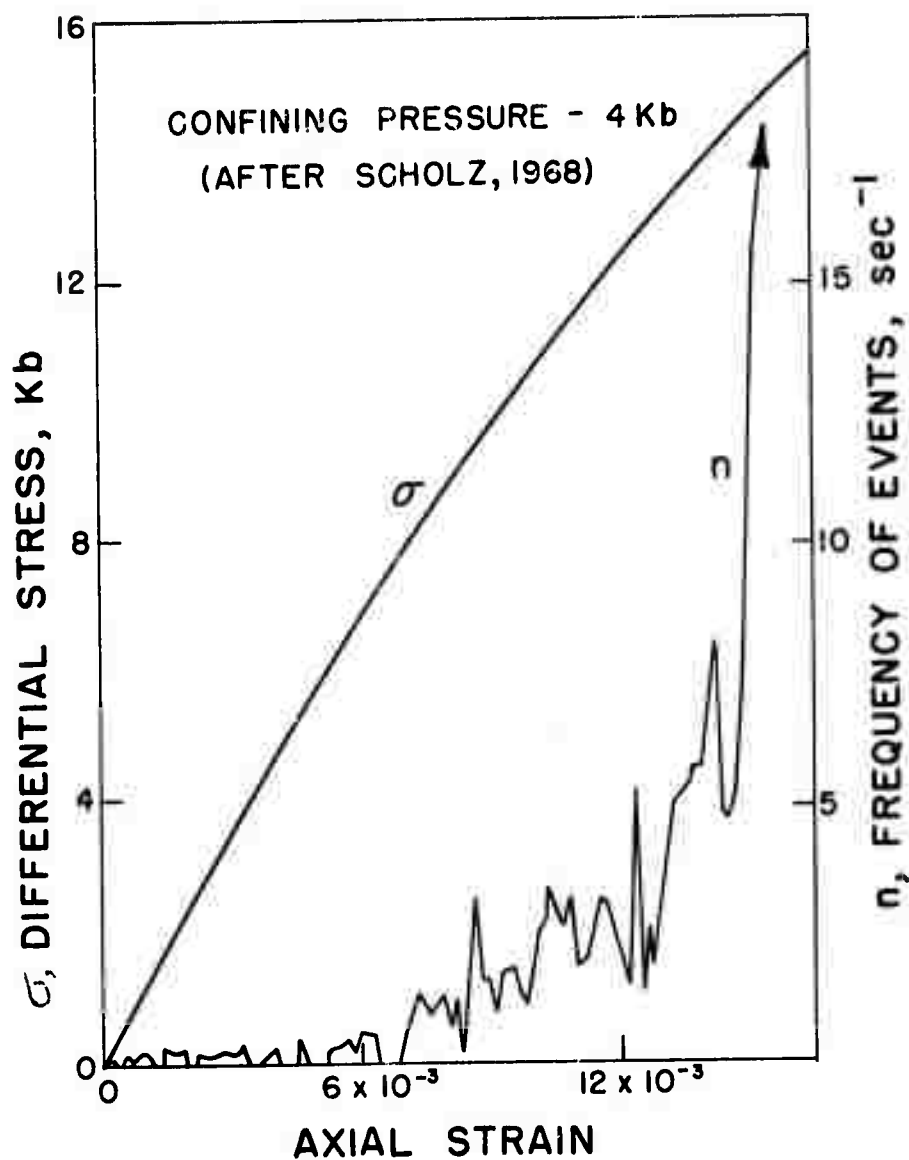
Figure 5. Microfracturing Frequency and Differential Stress Versus Strain for Westerly Granite Compressed Under 4 kb Confining Pressure. (After Scholz, 1968, Figure 6).

13
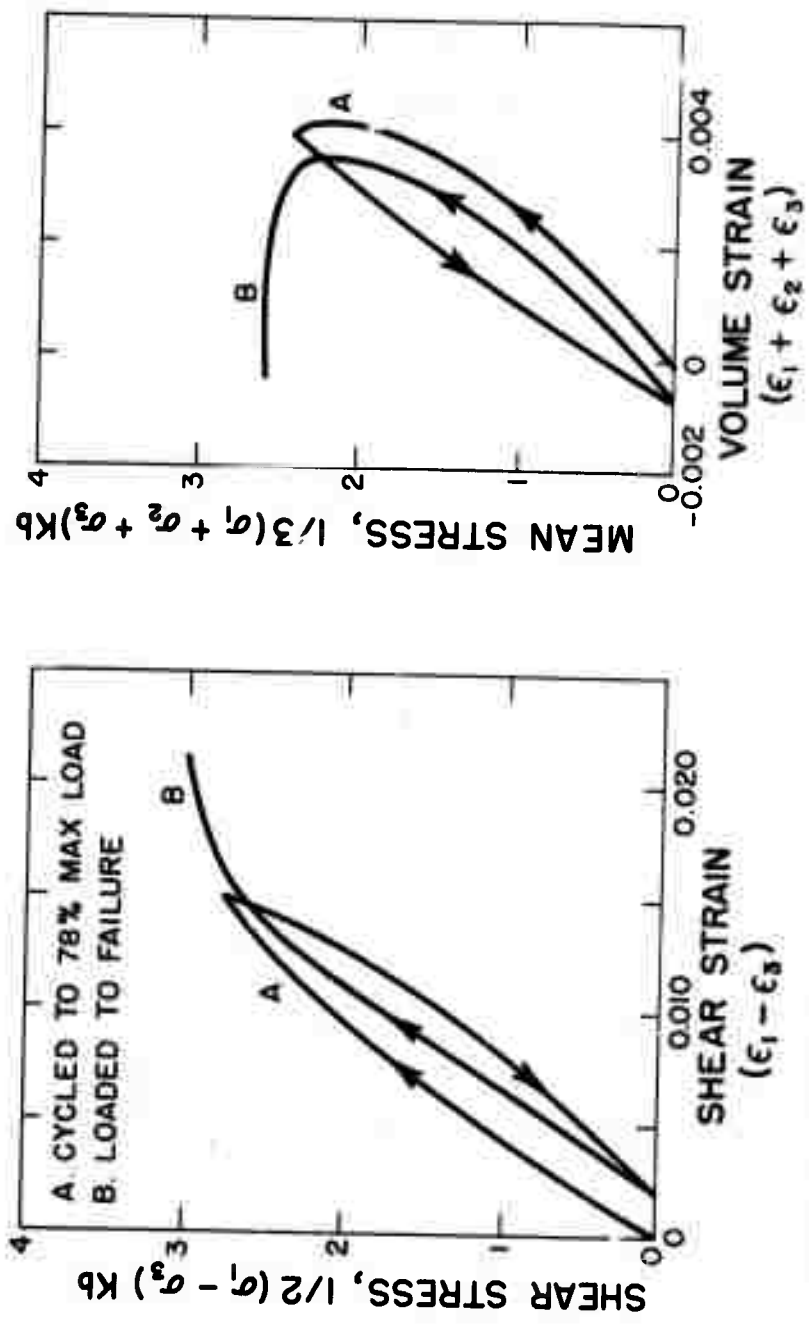
Figure 6. Shear Stress-Strain Curve and Volume Strain-Mean Stress Curve for Westerly Granite Under Cyclic, Proportional Loading with $\sigma_3/\sigma_1 \sim 0.1$. (After Swanson, 1969, Figures 32, 33).

14

The two important macroscopic defects in rock masses are layers and fractures (or joints). Both can impart a significant strength anisotropy to the medium. To assess the influence of these defects on the deformation we must compare the frictional properties of the interlayer and fracture surfaces with the strength properties of the intact rock.

If pre-existing fractures (or joints) do not have strong preferred orientations, they may merely degrade ultimate strength without imparting important directional properties to the rock. Figure 7 compares the stress-strain curves from triaxial-compression tests on intact Westerly granite with those obtained for the randomly precracked rock over about the same range of confining pressure. At low pressures the cracked rock is clearly much the weaker. With increasing pressure the disparity becomes smaller, until at 4.8 kb the intact rock, though somewhat the stiffer, is no stronger than the cracked rock. In other words, the higher the mean pressure, the less important are the defects, as one would expect.

If the fractures are preferentially oriented with respect to the stress field, they do impart a planar anisotropy to the rock which must be accounted for in our calculations of deformations.

The sliding friction of rock surfaces has been investigated by several workers during the past ten years or so. Because the dimensions of natural fractures and the spacings between them are generally large relative to the maximum specimen size available to most laboratories, systematic studies have been done largely on artificial surfaces created with a diamond saw.

Let's compare the resistance to sliding with the intact strength in triaxial-compression tests of 2 by 5-cm, jacketed cylinders of dolomite, sandstone, and limestone (Handin, 1969a). In the intact state these rocks are statistically homogeneous and isotropic in that breaking strengths are reproducible and independent of load orientation. Ultimate compressive strengths are nearly linear functions of confining pressure to at least 1 kb.

Specimens were cut at 5-deg increments from 30 deg to 70 deg to their longitudinal axes. Cuts of lower inclinations would intersect the ends. To facilitate correlations between tests, the surfaces were lightly polished merely to make them all as nearly alike as possible.

Figure 8 shows a typical force-time record of a triaxial test on a sample of sandstone cut at an angle θ_s of 45 deg to the maximum principal compression. Initially the constant force is due to the pressure of the confining fluid on the loading piston. When the piston contacts the specimen, the force rises linearly as the rock is compressed elastically. The force drops suddenly as slip on the cut begins and then oscillates because of stick-slip. Whether sliding occurs stably or by a stick-slip depends on the nature of the rock

15
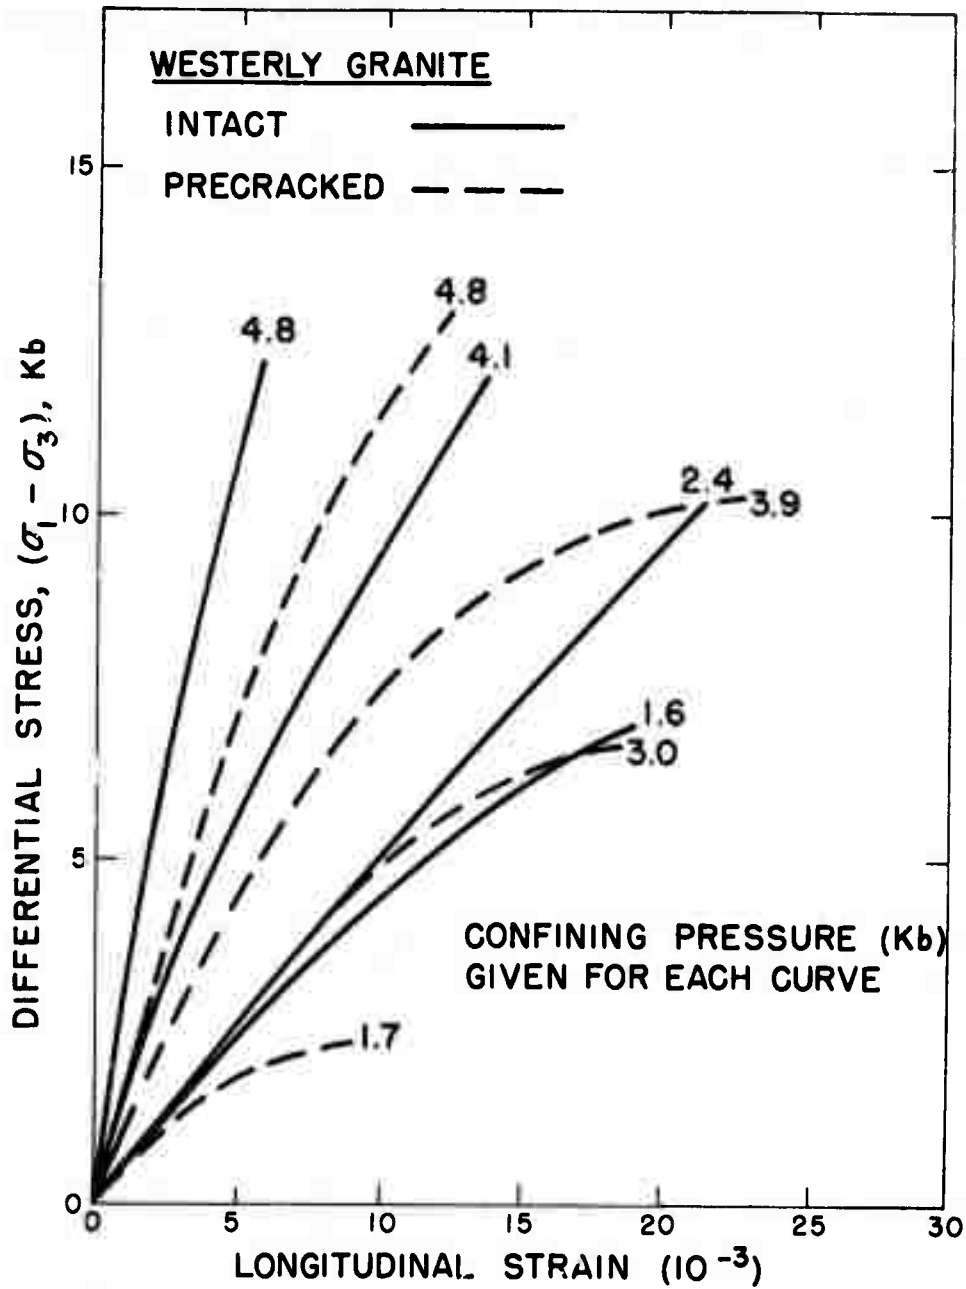
Figure 7. Stress-Strains for Intact and Precracked Westerly Granite Compressed Under Different Confining Pressures. (After Smith et al., 1969, Figures 5, 6).

16

surface, on normal stress, temperature, and probably strain rate, and on the inertial properties of the testing machine (Brace and Byerlee, 1966). The force drops suddenly to its initial value when the jacket breaks after large offset on the cut. The longitudinal maximum principal compression σ_1 at the onset of sliding is equal to the differential force ΔF divided by the cross-sectional area plus the confining pressure. The lateral minimum principal compression σ_3 is, of course, just equal to the confining pressure. The angle of the cut θ_s, or of a shear fracture θ_f, is measured clockwise from σ_1. The normal stress σ and shear stress τ on the cut are calculated from the equations in the lower left, and the coefficient of static friction μ is taken as the ratio of τ to σ. A separate run was made at each increment of confining pressure from 250 to 2000 bars because the process of slip itself may modify the nature of the surface.

$$\sigma = \frac{\sigma_1 + \sigma_3}{2} - \frac{\sigma_1 - \sigma_3}{2} \cos 2\theta_s$$

$$\tau = \frac{\sigma_1 - \sigma_3}{2} \sin 2\theta_s$$
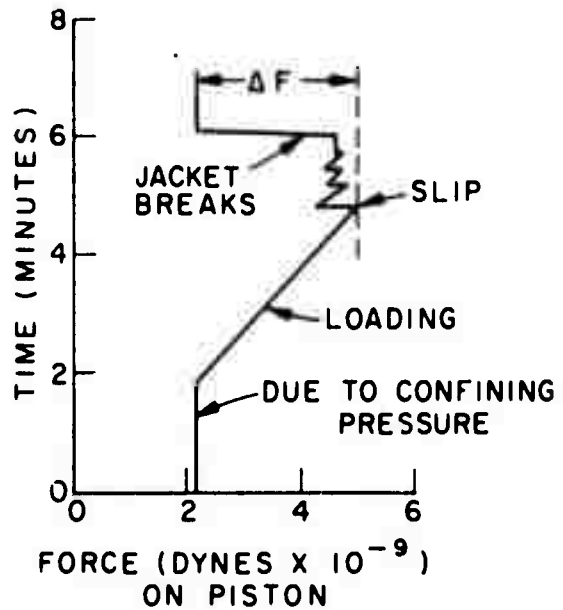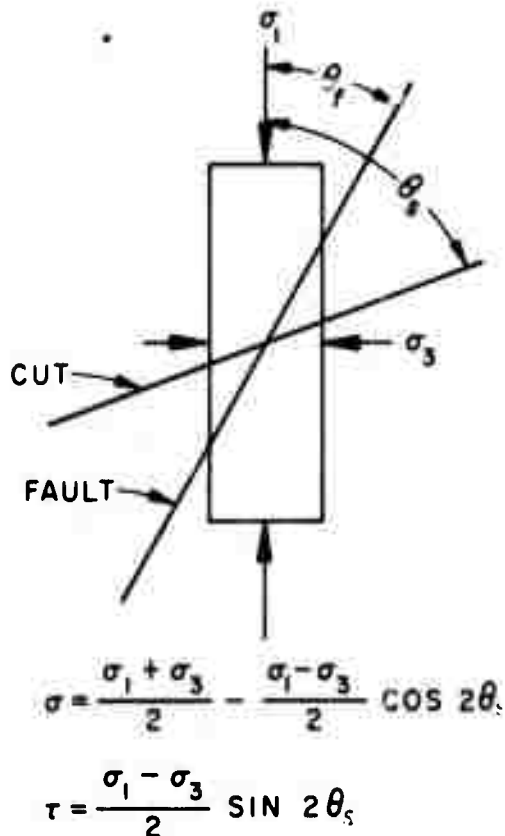
$\sigma_3 = \text{CONFINING PRESSURE}$

$\sigma_1 = \Delta F / \text{AREA} + \sigma_3$

Figure 8.   Stresses on Sliding Surfaces in Cylindrical Specimens Under Triaxial Compression (left). Typical Force-Time Curve for a Specimen Transected by a Saw Cut (right).

17

Figure 9 shows the coefficients of sliding friction, μ, on 45-deg saw cuts in four rocks as functions of normal stress, σ.

In Tennessee sandstone, the constituents of which are brittle under the test conditions imposed here, the decrease in μ is relatively small, 11 percent over the normal stress range of about 1 to 6 kb. Byerlee (1967a) has suggested that at low normal stresses surfaces slide by lifting over asperities, whereas at higher stresses the asperities must be broken through. If the material is in the brittle state, then in the latter stage the shear stress for sliding should increase linearly with normal stress as the Coulomb criterion predicts, that is, μ should remain constant. The behavior of this sandstone does not differ widely from this idealization, since the rate of change of μ is high at low normal stress and very small at high normal stress.

In the dolomites the decreases in μ with increasing normal stress from 0.5 to 3.5 kb are large, 27 percent for Blair and 22 percent for Knox. The changes of μ are about the same, and they are nearly constant. The only ready explanation for these reductions of the coefficients of friction is that these rocks are passing through the brittle-ductile transition and are behaving as do "plastic" metals in which a variable μ is well known. A large reduction (18 percent) is also observed in Solenhofen limestone, and here the rate of change of μ definitely increases above a normal stress of 4 kb. The corresponding mean pressure is about 3.5 kb, which is well above the 2.7-kb transition pressure.

To investigate the preference for shear fracturing or for slip on the cohesionless cut, we can superimpose a sliding line on the Mohr diagram (Figure 10). From a particular pair of values of the extreme principal stresses at failure, we construct a stress circle whose radius is half the stress difference and whose center lies on the normal-stress axis at half the sum of the principal stresses. The linear Mohr envelope, which represents the Coulomb-fracture condition, is tangent to the circle, has a slope equal to the tangent of the angle of internal friction ϕ, and intersects the shear-stress axis at τ_0, the cohesive shear strength. The tangent point gives the values of normal stress σ_f and shear stress τ_f on the fracture plane which lies at an angle of $2\theta_f$ from σ_1, as measured clockwise from the normal-stress axis. Fracture is supposed to occur when the total shearing resistance is equal to the cohesive strength plus the internal friction $\sigma_f \tan \phi$ on planes making an angle θ_f of 45 deg minus half the angle of internal friction. This criterion is not mechanistically satisfactory, but it does predict both breaking strength and fracture angle pretty well for the particular state of stress of these experiments.

The sliding line for a cohesionless cut passes through the origin and has a slope equal to angle of sliding friction, the angle whose tangent is the coefficient of sliding friction μ. The two pairs of coordinates at the intersections of the line and the Mohr stress circle satisfy the condition for sliding that the shearing resistance
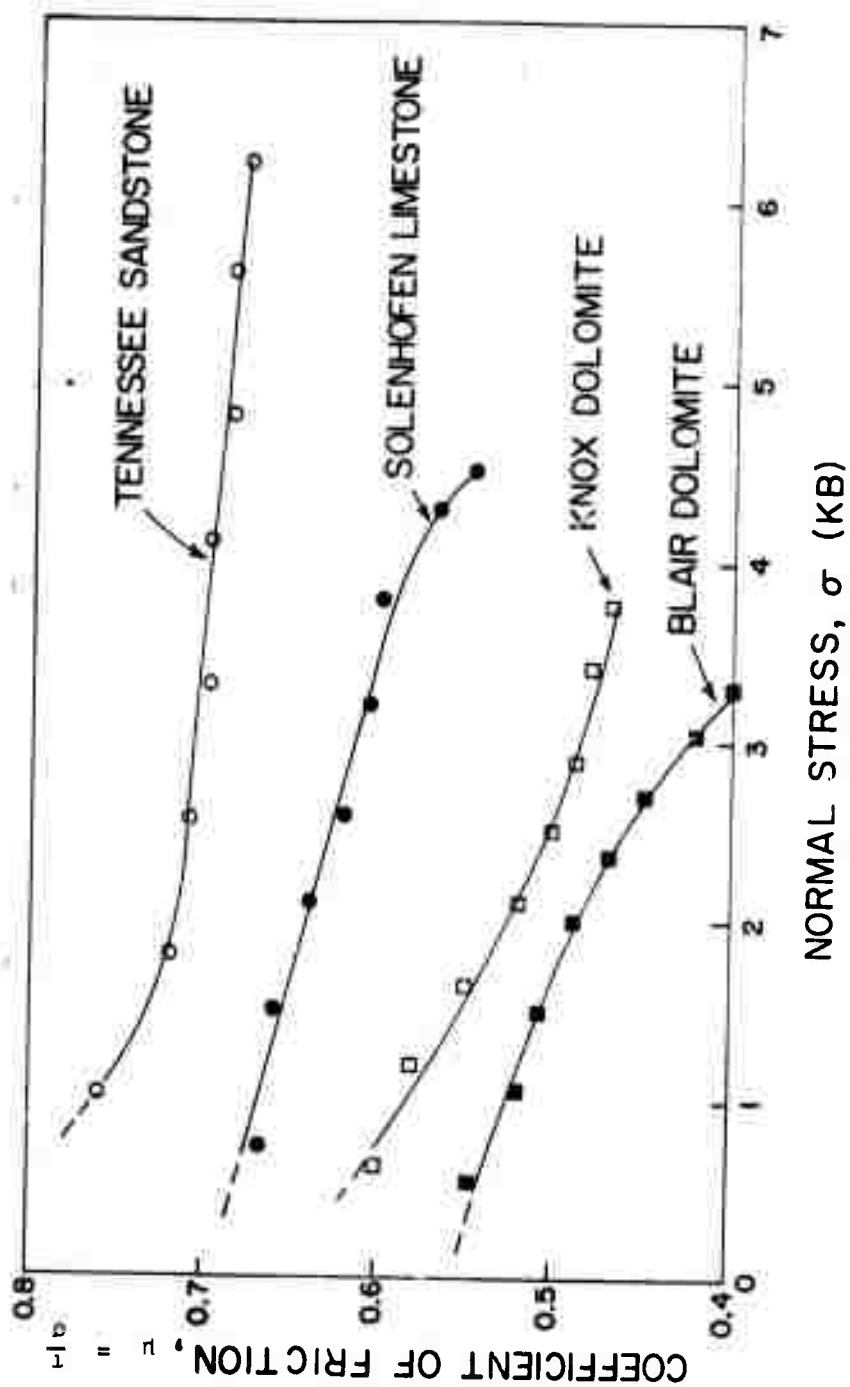
Figure 9. Coefficients of Sliding Friction on 45-deg Saw Cuts in Four Rocks as Functions of Normal Stress. (After Handin, 1969a, Figure 1).

19

τ_s equals the product of the coefficient of friction and the normal stress σ_s. That is to say, sliding on these two planes is just as likely as fracture, and sliding is favored on all planes lying between them. One of these is at a low angle to σ_1 and would intersect the ends of the specimen, thus it was not investigated. To simplify already cluttered diagrams, sliding lines will be extrapolated linearly through the origin. Slip on the other plane, inclined at an angle θ_s of 30 deg or more to σ_1, can be compared with fracture.
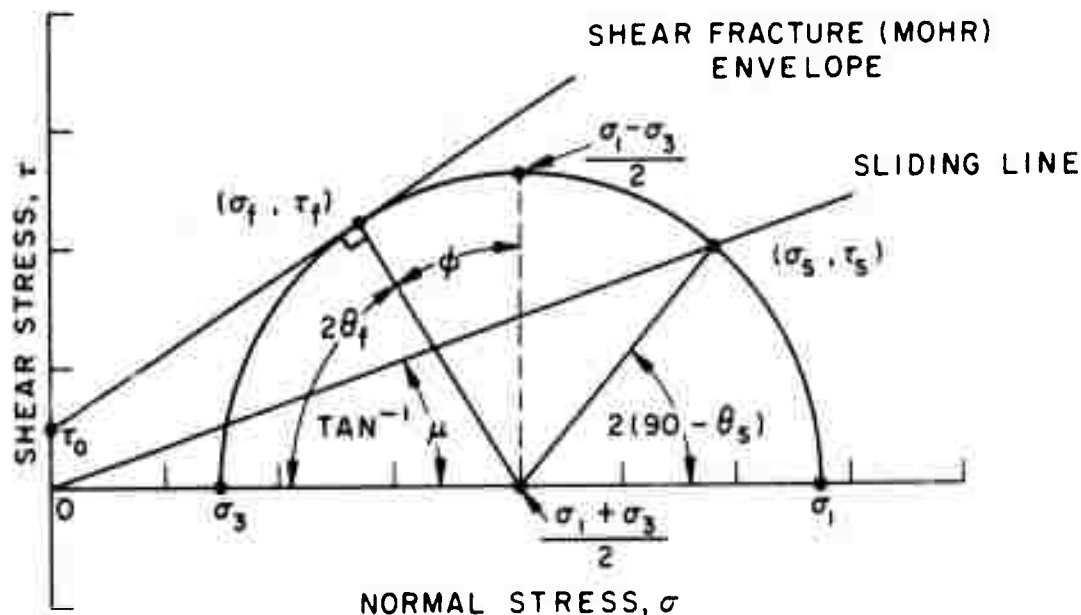
COULOMB FRACTURE

$$\tau_f = \tau_0 + \sigma_f \; TAN \; \phi$$

$$\theta_f = \pm 45° \mp \frac{\phi}{2}$$

SLIDING

$$\tau_s = \mu \, \sigma_s$$

Figure 10. Mohr Diagram with Shear Fracture Envelope and Line for Sliding on a Cohesionless Cut.

Consider some test data at a confining pressure, σ_3 of 1 kb (Figure 11). Consider Blair dolomite with cohesive strength of 450 bars and angle ϕ of 45 deg. For the normal stresses on high-angle cuts, μ is about 0.4 at this confining pressure. The contrast between the angle of internal friction ϕ, 45 deg, and the angle of sliding friction, 21 deg, is the largest encountered. We predict that slip on a 65-deg cut is as likely as fracture. Indeed we observe that sliding alone always occurs at 60 deg and less. Fracturing alone occurs in 70-deg cylinders at about 20 deg to σ_1 in accord with the Coulomb condition.
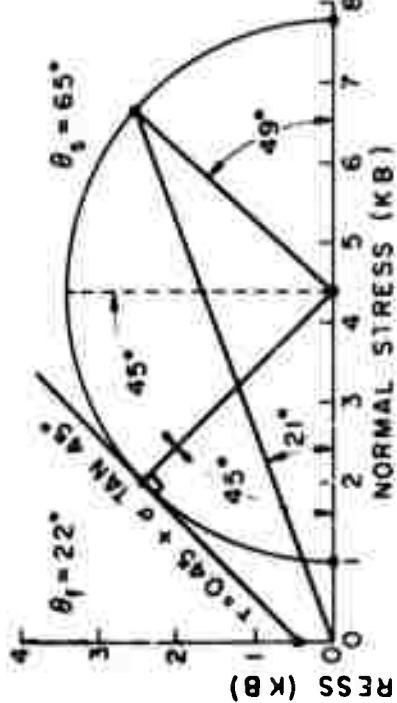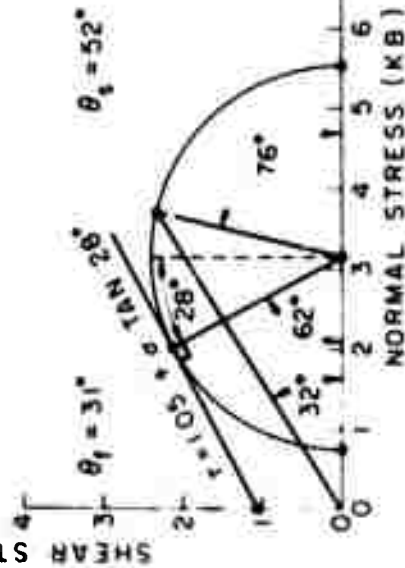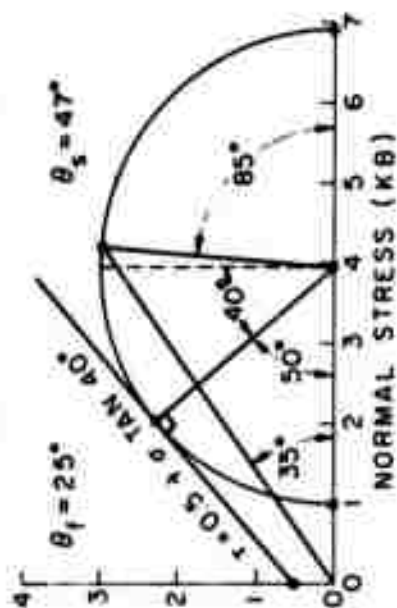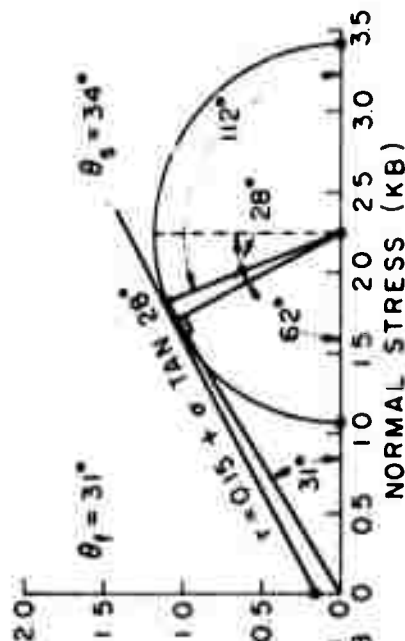
Figure 11. Mohr Envelopes and Sliding Lines for Four Rocks. (After Handin, 1969a, Figure 2).

21
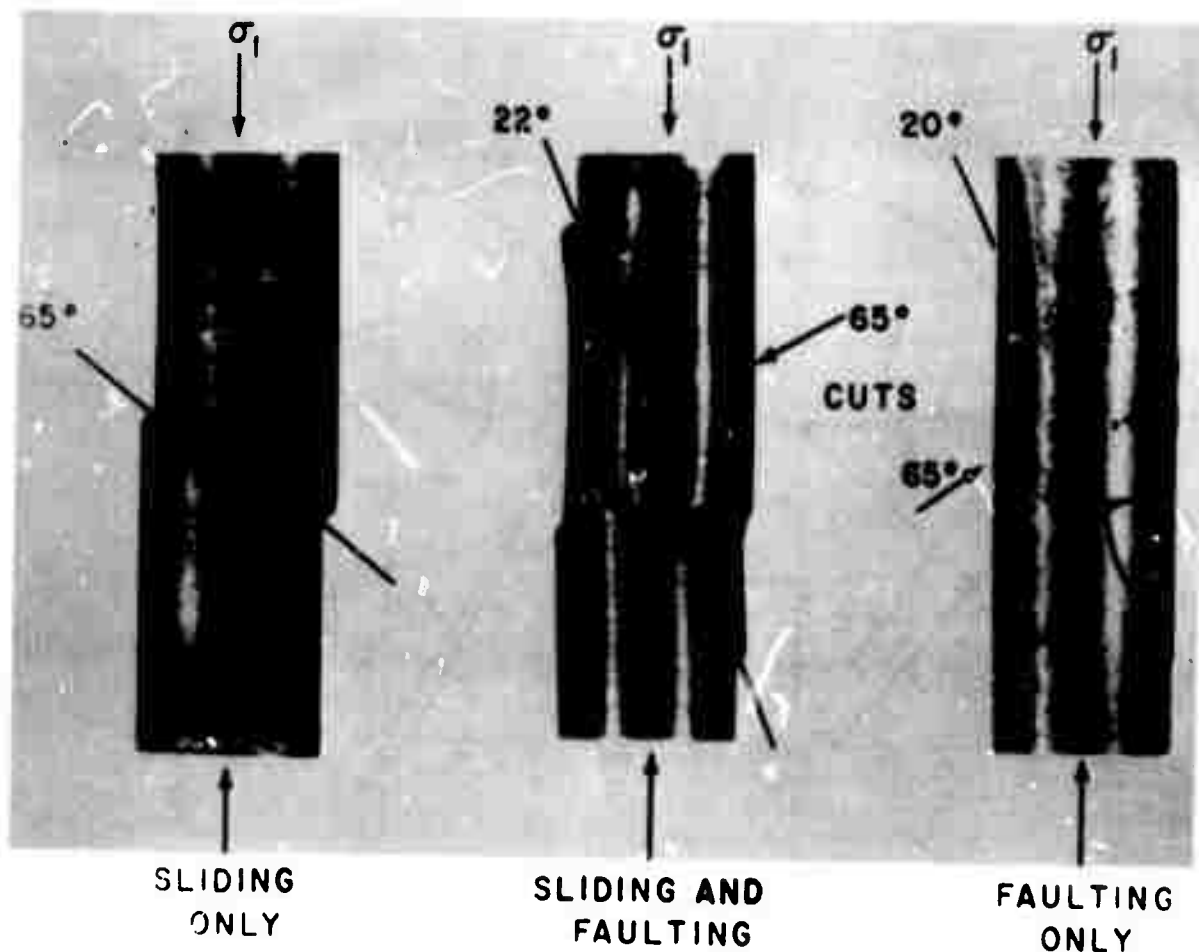
Figure 12.   Blair Dolomite Specimens with 65-deg Saw Cuts.
(After Handin, 1969a, Figure 3).

Either sliding, fracturing, or both is seen in specimens with
65-deg cuts (Figure 12). In the central cylinder the fault has offset
the cut, but to a lesser degree, sliding has offset the fault as well.
The traces of the cut and the fracture on the circular section are
nearly always parallel, but dip directions are inconsistent.

Tennessee sandstone has a τ_0 of 500 bars, ϕ of 40 deg, and the
highest angle of sliding friction, about 35 deg ($\mu = 0.7$) (Figure 11).
We predict that fracture at 25 deg is preferred for cuts of 50 deg or
more, and this is just what we see.

For Solenhofen limestone the angle of sliding friction, 32 deg
($\mu = 0.62$), exceeds that of internal friction, 28 deg, and the 1050-bar
cohesive strength is very high, so that fracturing at 31 deg rather than
sliding should occur at 55 deg and more.

For Lueders limestone the 31-deg sliding friction angle ($\mu = 0.6$) also exceeds the internal friction angle of 28 deg. Here the cohesive strength is so low, only 150 bars, that the sliding line almost fails to intersect the stress circle at all. Slip should become barely possible at 34 deg. Displacement on a 30-deg cut does occur, but since this surface lies very near the 31-deg plane of potential failure, sliding and faulting are no longer really distinguishable.

If the mean pressure had been a little higher, the limestone would have deformed as if it had been intact without regard to the pre-existing joint. The Mohr envelopes and sliding lines of all rocks eventually curve toward the normal-stress axis if only the mean pressure is high enough. Once these two curves intersect, sliding is no longer possible, and the effects of jointing can be ignored--that is, the rock can be treated as if it were intact. For strong rocks like granite this condition will probably be reached at effective confining pressures somewhere between 12 and 20 kb. More work is needed on this problem.

What is the effect of pore pressure on sliding friction? By plotting the shear stress against the effective normal stress for the maximum friction on ground surfaces of water-saturated Westerly granite, Byerlee (1967b) has shown that the effective-stress concept holds (Figure 13). The slopes of the curves for three different pore pressures are all the same, namely, 0.6. This result implies, of course, that in a saturated, jointed rock mass, block motion will probably occur no matter how high the total stresses may be. Because the framework compressibility may well be less than that of water, rapid loading will tend to raise pore pressure and to reduce the shear strengths of joints.

The measurements on artificial surfaces yield coefficients of friction in the range of 0.4 to 0.8 for the common rock types. Are these values valid for natural fractures? The few tests done to date suggest that they are in fact representative of the real rock mass. The Corps of Engineers (1965) has compared the Mohr envelope of intact gneiss with those of samples containing dry, open natural joints, dry saw cuts, and wet saw cuts (Figure 14). Note that natural and artificial joints have the same sliding line. Also note that the shear strength of the joint in the water-saturaged gneiss is nearly vanishing.

The dramatic effect of layering on strength has been demonstrated by Donath (1961) in triaxial compression tests on a slate with a pronounced planar anisotropy due to cleavage (Figure 15). From a series of tests at three different confining pressures he has plotted ultimate compressive strength against the inclination of the cleavage with respect to the maximum principal compressive stress, σ_1. The highest strengths are measured at 0 deg to 90 deg--that is, parallel and normal to cleavage. Thirty-degree specimens have the lowest strength. These results imply that the cohesive strength, the internal friction, or both are lower in the cleavage plane than elsewhere in the

23
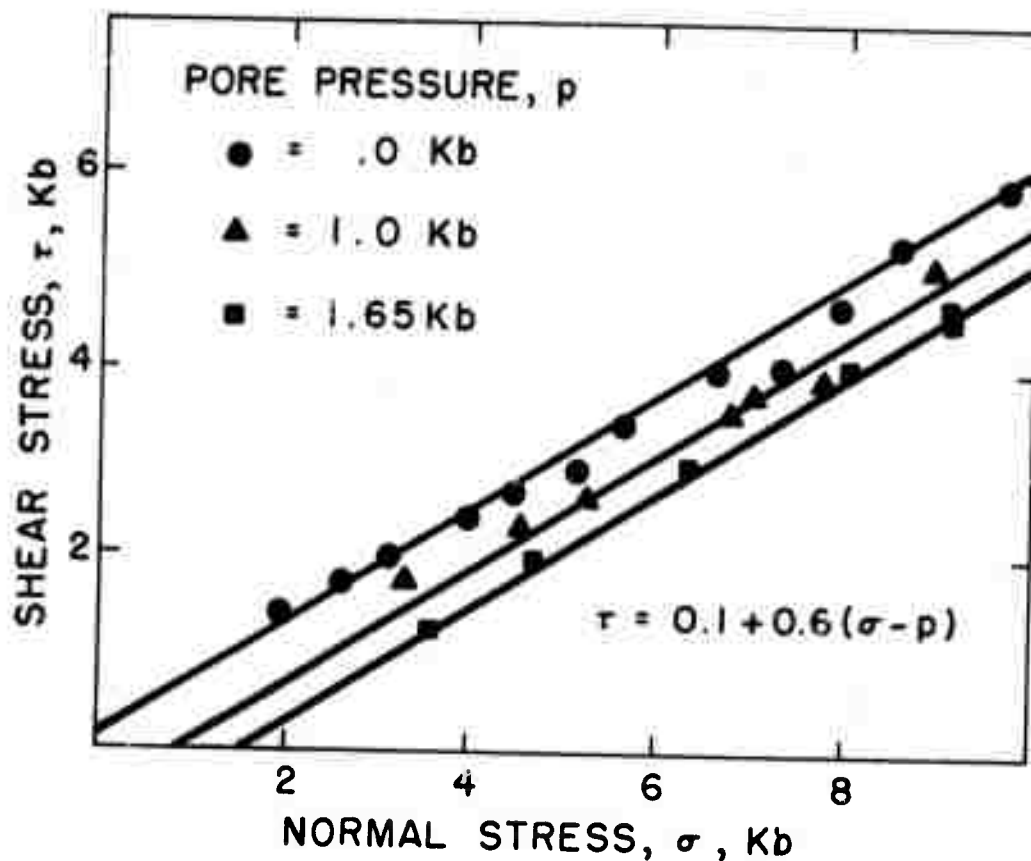
Figure 13. Shear Stress Versus Normal Stress for Maximum Friction
on Ground Saw Cuts in Water-Saturated Westerly Granite.
(After Byerlee, 1967, Figure 5).

24
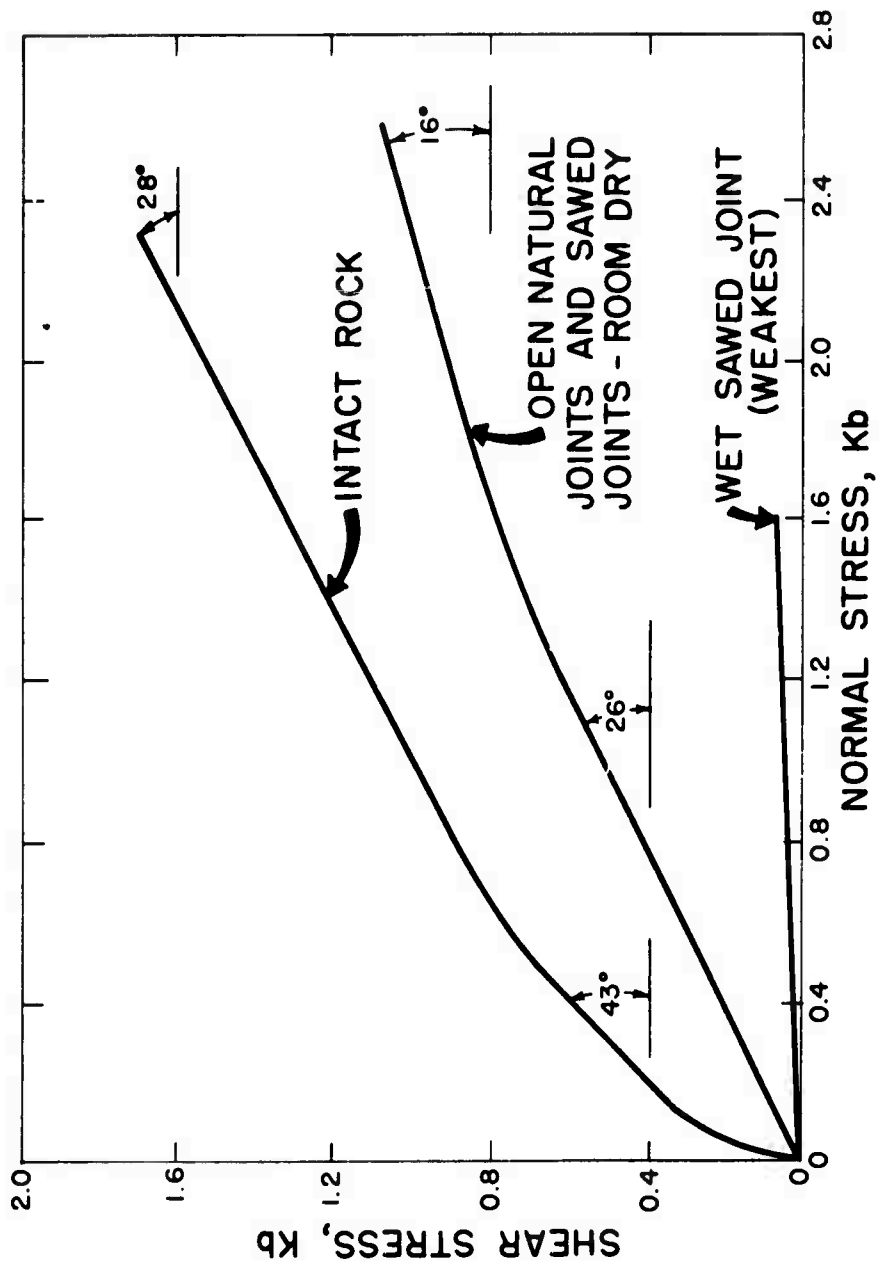
Figure 14. Mohr Shear-Fracture Envelope and Sliding Lines for Natural and Artificial Surfaces in Dry and Water-Saturated Schistose Gneiss. (After Corps of Engineers, 1965, Figure 3).

25
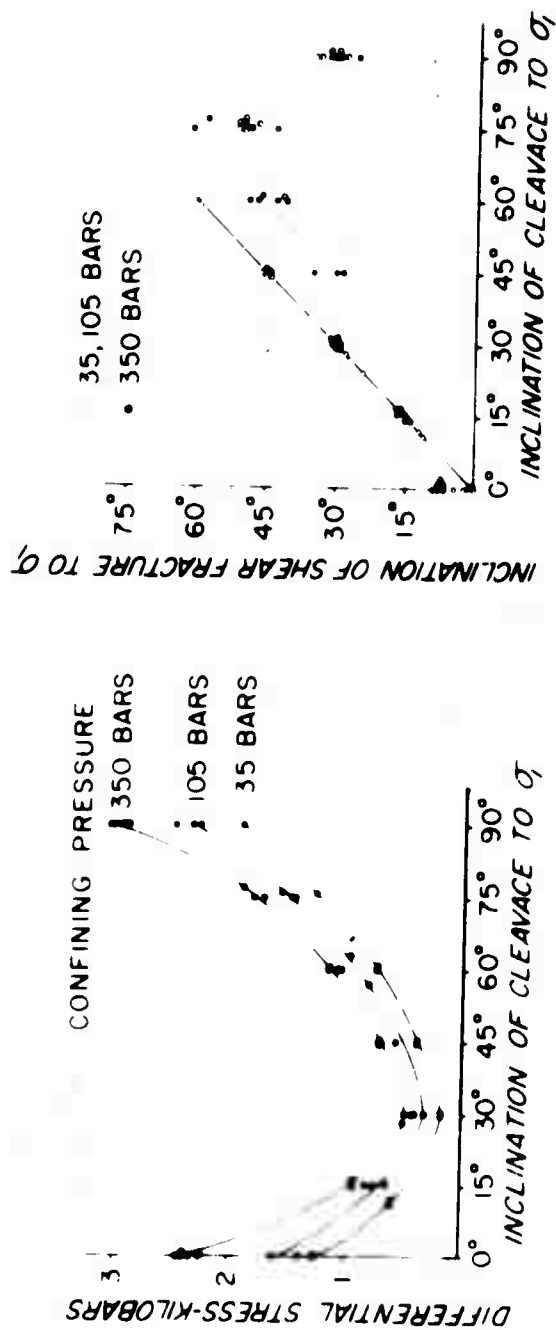
Figure 15. Effects of Cleavage on the Fault Angle and Ultimate Compressive Strength of Martinsburg Slate at Different Confining Pressures. (After Donath, 1961, Figures 1, 2).

26

rock. We do not know to what mean pressures this anisotropy will persist. We do know from other work that gneiss, schist, and slate are still strongly anisotropic under 5 kb confining pressure (Borg and Handin, 1966).

Cleavage also influences the orientations of shear fractures. For inclinations of cleavage from about 5 deg to 45 deg the fractures tend to follow the layering. In 90-deg specimens fracturing occurs at about 30 deg to σ_1 as usual. In 60-deg and 75-deg specimens, although fracturing does not follow cleavage, it is influenced by the anisotropy as it occurs at about 45 deg and not 30 deg to σ_1.

In summary, laboratory techniques for measuring the elastic properties and static-strength properties are well established, and the phenomenological behavior of intact homogeneous rocks is pretty well understood. Further experimental work is needed on anisotropic rocks with layering, jointing, or both. We need more dynamic testing of both intact and cracked material--both wet and dry--over the entire range of confining pressure from 1 bar to 100 kb.

It should be obvious that because of the heterogeniety of rock, its properties as associated with any particular ground-motion prediction should be determined on samples collected at or near the shot point.

We knew much of this at our last meeting 2-1/2 years ago. What has been accomplished in the meantime? I think there have been several significant advances:

(1) We are beginning to collect gas-gun data on cracked, water-saturated rocks up to about 100 kb and reliably to map release paths (e.g., Lysne, 1970; Peterson et al., 1969).

(2) We have improved our capability to deform large samples and hence to test natural defects under moderate confining pressure. I should add that we have not always utilized already available facilities. The Bureau of Reclamation laboratory can handle 15 by 30-cm cores up to 8 kb.

(3) Because much ground motion occurs during unloading, it is important to obtain complete stress-strain curves. Using very stiff testing machines, Wawersik (1968), Rummel (1970), and others have obtained a few such curves, exemplified in Figure 16. In very brittle rocks like the basalt and Solenhofen limestone, fracture still occurs too rapidly for delineation of the entire unloading path. Modern electrohydraulic, servo-controlled testing machines of high apparent stiffness may yield more nearly complete curves of brittle materials than could be measured previously. In any event, curves are now much easier to obtain.

27
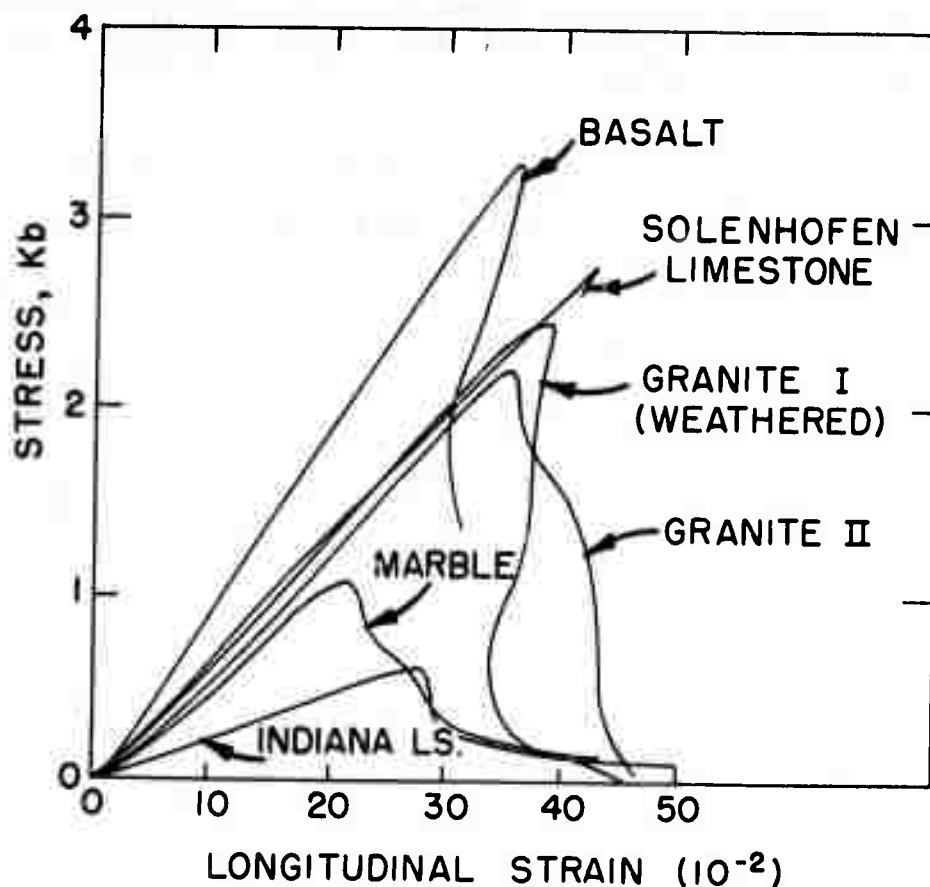
Figure 16. Stress-Strain Curves for Rocks from Uniaxial Compression
Tests in a Stiff Machine. (After Wawersik, 1968, Figure 3).

(4) These machines can also be programmed for one-dimensional compres-
sion tests, proportional-loading tests, and just about any dynamic
loading up to 10 Hz. W. F. Brace will comment later on some of his
recent work; also see Swanson (1969).

(5) We have begun to measure the uniaxial strength properties of intact
rocks at intermediate strain rates up to about 10^4 per second
(Green and Perkins, 1970) and under confining pressures to 8 kb
at about 10 per second (Handin, 1969b). We are getting better
measurements of Hugoniot elastic limits at very high strain rates
and mean pressures. All this work will soon be extended to
cracked, water-saturated rocks.

(6) Several investigators are working with block models of jointed rock
under two-dimensional compression, and they are comparing the
results with finite-element analyses. Quantitative measurements
of the response of rock to explosive loading are underway
(Godfrey, 1969). Unfortunately, it would be rare and fortuitous

if the symmetry of a natural fracture pattern with respect to explosive loading allowed for less than full three-dimensional analysis.

As I see it, the most challenging need now is for cheap, reliable methods of in situ measurements, especially of strength, under both static and dynamic loading. To be significantly advantageous over laboratory testing, the field testing must sample a large volume, preferably several cubic meters, without disturbing the rock mass. The ideal method would be nondestructive. Attempts have been made, for example, to correlate seismic velocity, attenuation, or both with mechanical properties, but so far the results have been only qualitative.

Another closely related problem is the measurement of the in situ state of stress. The imprints of tectonic prestress are now clear on both the teleseismic radiation pattern and the close-in surface strains associated with many large underground explosions. It is also clear that seismogenic faulting has been triggered. Although the largest magnitude of aftershocks has been two orders of magnitude below that of the explosion itself (Evernden, 1969), and the most distant epicenter has been only 40 km from the shot point (Boucher et al., 1969), this phenomenon also compounds the difficulty of the seismic-detection problem. There is need for stress measurements in deep boreholes, and some techniques are now under development.

I left our last meeting a couple of years ago with the impression that each group of people working on rock properties, code calculations, and field tests might as well have met in three separate rooms. In Albuquerque two years later the situation seemed no better.

There have been too many instances of indiscriminate culling of the rock-properties literature, without regard to experimental procedures, precision of data, and the tested rock as a valid analog of the medium. It's up to rock-properties people to provide reliable, pertinent data. It's up to the code people to check the sensitivity of their calculations and to determine which parameters to measure and how precisely. And it's up to the test people to provide adequate descriptions of the media and unassailable records of the real ground motions for comparisons with predictions.

I don't know why we have been unable to communicate with each other more effectively. I do know that until we learn how, we can't solve the problems despite all the enormous talent available. Let's hope this meeting brings us closer together.

References Cited

Borg, I., and J. Handin, Experimental deformation of crystalline rocks, *Tectonophysics*, v. 3, p. 249-368, 1966.

Boucher, G., A. Ryall, and A. E. Jones, Earthquakes associated with underground nuclear explosions, *J. Geophys. Research*, v. 74, p. 3808-3820, 1969.

Brace, W. F., The mechanical effects of pore pressure on the fracturing of rocks, Research in Tectonics, Geol. Surv. Canada, *Paper 68-52*, 1968.

Brace, W. F., and J. D. Byerlee, Stick-slip as a mechanism for earthquakes, *Science*, v. 153, p. 990-992, 1966.

Brace, W. F., B. W. Paulding, Jr., and C. Scholz, Dilatancy in the fracture of crystalline rocks, *J. Geophys. Research*, v. 71, p. 3939-3953, 1966.

Byerlee, J. D., Theory of friction based on brittle fracture, *J. Appl. Physics*, v. 38, p. 2928-2934, 1967a.

Byerlee, J. D., Frictional c acteristics of granite under high confining pressure, *J. Geophys. Research*, v. 72, p. 3639-3648, 1967b.

Corps of Engineers, Tests for strength characteristics of a schistose gneiss, Missouri River Div. Lab., *Rept. 64-126*, 20 p, 1965.

Donath, F. A., Experimental study of shear failure in anisotropic rocks, *Geol. Soc. America Bull.*, v. 72, p. 985-990, 1961.

Evernden, J., Magnitude of nuclear explosions and their aftershocks (Abstract), *Am. Geophys. Union Trans.*, v. 50, p. 247, 1969.

Godfrey, C. S., and H. Y. Moises, A dilatant model for the response of rock to buried explosions, Physics International Co., *Rept. PITR-69-3*, 10 p, 1969.

Green, S. J., and R. D. Perkins, Uniaxial compression tests at strain rates from 10^{-4} to 10^{4} per second on three geologic materials, *Proc. 10th Symposium on Rock Mechanics*, AIME, in press, 1970.

Handin, J., On the Coulomb-Mohr failure criterion, *J. Geophys. Research*, v. 74, p. 5343-5348, 1969a.

Handin, J., Studies in rock fracture, *7th Quart. Tech. Rept.*, Contract DACA73-68-C-0004, Office of the Chief of Engineers, Department of the Army, Washington, D.C., 1969b.

Handin, J., R. V. Hager, Jr., M. Friedman, and J. N. Feather, Experimental deformation of sedimentary rocks under confining pressure: Pore pressure tests, *Am. Assoc. Petroleum Geologists Bull.*, v. 47, p. 717-755, 1963.

Lysne, P. C., A comparison of calculated and measured low stress Hugoniots and release adiabats of dry and water-saturated tuff, *J. Geophys. Research*, v. 75, p. 4375-4386, 1970.

Petersen, C. F., W. J. Murri, and R. W. Gates, Dynamic properties of rocks, Stanford Research Institute, *Rept. PGU-6273*, 46 p, 1969.

Rummel, F., and C. Fairhurst, Determination of the post-failure behavior of brittle rock using a servo-controlled testing machine, *Rock Mechanics*, in press, 1970.

Scholz, C. H., Microfracturing and the inelastic deformation of rock in compression, *J. Geophys. Research*, v. 73, p. 1417-1432, 1968.

Smith, J. L., K. L. De Vries, D. J. Bushnell, and W. S. Brown. Fracture data and stress-strain behavior of rocks in triaxial compression, *Experimental Mechanics*, v. 9, no. 8, p. 348-355, 1969.

Swanson, S. R., Development of constitutive equations for rocks, Ph.D. Thesis, Univ. Utah, 140 p, 1969.

Terzaghi, K., *Theoretical Soil Mechanics*, New York, Wiley and Sons, 510 p, 1943.

Wawersik, W. R., Detailed analysis of rock failure in laboratory compression tests, Ph.D. Thesis, Univ. Minnesota, 165 p, 1968.

CHAIRMAN SIMMONS: Thank you John. I am sure that all of us realize the amount of effort needed to prepare an overview of such a large field.

I would like to turn now to our panel members and ask for comments, both specific and general. Some of you have come prepared with slides and I think you should let me know when it is appropriate to mesh in your particular five or ten minute comment. The floor is open for discussion by the panel members.

MR. CHERRY: I would like to make one comment first and then ask a question about failure criteria in general.

It seems to me that a failure criterion ought to include two things: It ought to provide a statement of when failure occurs in the rock, and it also ought to include a description of the stress adjustments that need to be made while failure is going on. So that is the statement.

The question I would like to ask is "Do the current (more or less) standard rock mechanic strength tests that we have available to us, like triaxial compression and extension and so forth, have any hope at all of describing the stress adjustments that need to be made during the process of unstable fracture propagation?" It seems to me that we are fairly close to getting some good, or we have some good, experimental data now as far as the region of stable fracture property is concerned. What about the unstable fracture propagation? What happens there? Are we really ignoring that phenomenon in the experimental efforts that are going on?

CHAIRMAN SIMMONS: Do you have a feeling for that, John?

Let me say, also, that I don't look to our speakers to answer every question that comes up here.

MR. CHERRY: I don't know who I am addressing it to. I would like the experimentalists to help answer it.

CHAIRMAN SIMMONS: Do you want to take a hack at that, John?

MR. HANDIN: I will take a hack at it. I think Bill Brace, who has really been more concerned with the details of fracture propagation than I have, also might wish to comment.

I don't think we are ignoring this problem because we don't think it is important, but I don't know how to measure the stress except to measure the average stress and consider the specimen as a whole. And what you want to know, I presume, is the present distribution of stress within the material after it failed.

Preceding page blank

MR. CHERRY: Yes. The only data I have seen so far are those of Byerlee, who tried to measure the stress drop after failure.

MR. HANDIN: Perhaps we are talking about two different things. I have already said it is important to get the complete stress-strain curves. In that sense, Byerlee is measuring sliding friction. I think we are on the verge of being able to get a complete stress-strain curve within the material. But this is still an average curve obtained by measuring the stresses on the boundaries of the specimen and really not internally. I don't know how you would measure the internal redistribution stress.

MR. CHERRY: The stress-strain curves that I have seen go up to the unstable fracture propagation. Then they stop; there are no further data.

MR. GODFREY: I was going to comment on your last slide. It seems to me that the particular unloading data typical of that kind of curve are not very useful because the boundary conditions on the jacket are completely unknown and uncontrollable. I don't think that that kind of unloading relates to reality.

MR. HANDIN: What would you like to see?

MR. GODFREY: I would like to see unloading data in a region where I can find the boundary conditions.

MR. BRACE: I think there is a lot of work these days on following the stress-strain behavior beyond the peak.

In particular, such work is being done by Byerlee, by the group at Utah, and by me. The conditions, though, are fairly restricted. For example, one keeps confining pressure constant beyond the peak and follows either the sliding or the subsequent violent or nonviolent motions that occur on the fractures. That is one very restricted situation, namely, that the external pressure on the sample remains the same.

There are also studies of the growth of individual cracks in both tension and compression in glass and synthetics. I assume they would be of interest here.

MR. GODFREY: I believe that the unloading conditions most useful to the codes are one-dimensional strains.

MR. BRACE: So one-dimensional, unloaded experiments would be the ones most useful. To my knowledge they have not been done, although the group at Utah provided some information.

CHAIRMAN SIMMONS: Why do you desire one dimensional?

MR. CHERRY: Because you can model it easily.

CHAIRMAN SIMMONS: Does this condition match the boundary conditions in the earth, though?

34

MR. CHERRY: If you have a one-dimensional model, and you do the corresponding test on it, that ought to be the first test that you do in order to check the model. Then you go to more complicated stress states. My computation budget is not unlimited.

MR. GODFREY: The cracks that we are talking about are small, certainly at distances where the seismic response can be considered almost plane waves and the lateral strains are very small. But I believe that once a small specimen starts failing then the effective cross-section area becomes almost indefinable.

CHAIRMAN SIMMONS: What about the boundary conditions on your small specimen?

Bill Brace, would you comment on whether the boundary conditions of the small laboratory specimen and pressure vessel match the boundary conditions on a similar volume of rock in the area under failure?

MR. BRACE: The kind of boundary conditions we can impose most easily in the laboratory are those of stress. What sort of boundary conditions do you think are appropriate?

I would like to turn it back to the other gentleman.

CHAIRMAN SIMMONS: I don't know, I am raising the question whether they are realistic for the earth.

MR. TRULIO: The important, simple one is uniaxial strength.

MR. BRACE: This condition can be simulated. Perhaps now might be the appropriate time to give my ten-minute talk on the uniaxial-strain experiments currently underway.

It is possible to subject material to uniaxial-strain boundary conditions in a static laboratory experiment. This sort of thing is done by ourselves, by the group consisting of Wayne Brown and Steve Swanson, and by others. In fact, the Utah group has pioneered this investigation.

Briefly it consists of working with a small sample on which strain gages are fixed. One applies an axial stress and a radial stress. One measures axial strain and circumferential strain, keeping the two loads, the axial and the radial, at such a ratio that the circumferential strain is zero. All the distortion is in the axial direction of the three principal strains and then only one strain is nonzero, hence, uniaxial strain.

There are two things that one observes in the laboratory experiment. One is the ratio of the two stresses required to maintain

the one strain zero, and the other is the volume change. So we get both
a stress-strain relation and a stress-volume change relation.

Figure 17 shows some of the details of the sample. Basically,
we put the strain gage, or several strain gages, on short cylindrical
samples. The radial stress is applied by fluid pressure acting on the
outside, and the sample is loaded axially.

Figure 18A shows results for a well-known rock, the Westerly
granite. I show two curves here*, plots of the two stresses, the radial
and the axial. The axial stress is plotted on the ordinate, the radial
stress, called the confining pressure, is plotted on the other axis.
The upper curve has been drawn through points at which this rock frac-
tures in a typical triaxial experiment. The lower curve is that which
one obtains from a uniaxial-strain experiment.

Several things are of interest. First, note that the uniaxial
strain falls well below the fracture curve. It doesn't seem to be
trending toward the fracture curve at high pressures. In other words,
it is still converging somewhat at the upper right hand part of the dia-
gram. Note that the stresses obtained here are considerable. The
stress is in excess of 35 kb required for a fracture. The stress in the
uniaxial case reaches 24 to 25 kb.

Let me describe quickly three or four results of this work
for the class of rock that includes granite.

The uniaxial stresses restricted to the uniaxial-strain con-
dition produce no observable yield and no observable cracking although
the axial stresses reach in some cases about 30 kb with 10 kb confining
pressure.

The second point is the ratio itself of these two stresses in
the uniaxial case. If this material were perfectly elastic, they would
be related to Poisson's ratio, and we could take points and compare with
what one gets in an ultrasonic measurement of the shear velocity at the
same confining pressure. The values of Poisson's ratio that you obtain
by taking this ratio and the one you take from the ultrasonic data, are
quite different.

For example, at the upper right hand end of this curve ultra-
sonics gives about 0.27 or 0.28 for Poisson's ratio. The values taken
from this curve are 0.33 to 0.35, so if one were to turn around and try
to calculate this there would be an appreciable error. The reasons for
this difference are rather intriguing.

The third point: One can compare stress-volume change from a
uniaxial experiment with pressure-volume change from shock experiments.

---

*Editor's Note: One of these two curves is missing from Figure 18A.

(a) Uniaxial

(b) Hydrostatic + triaxial
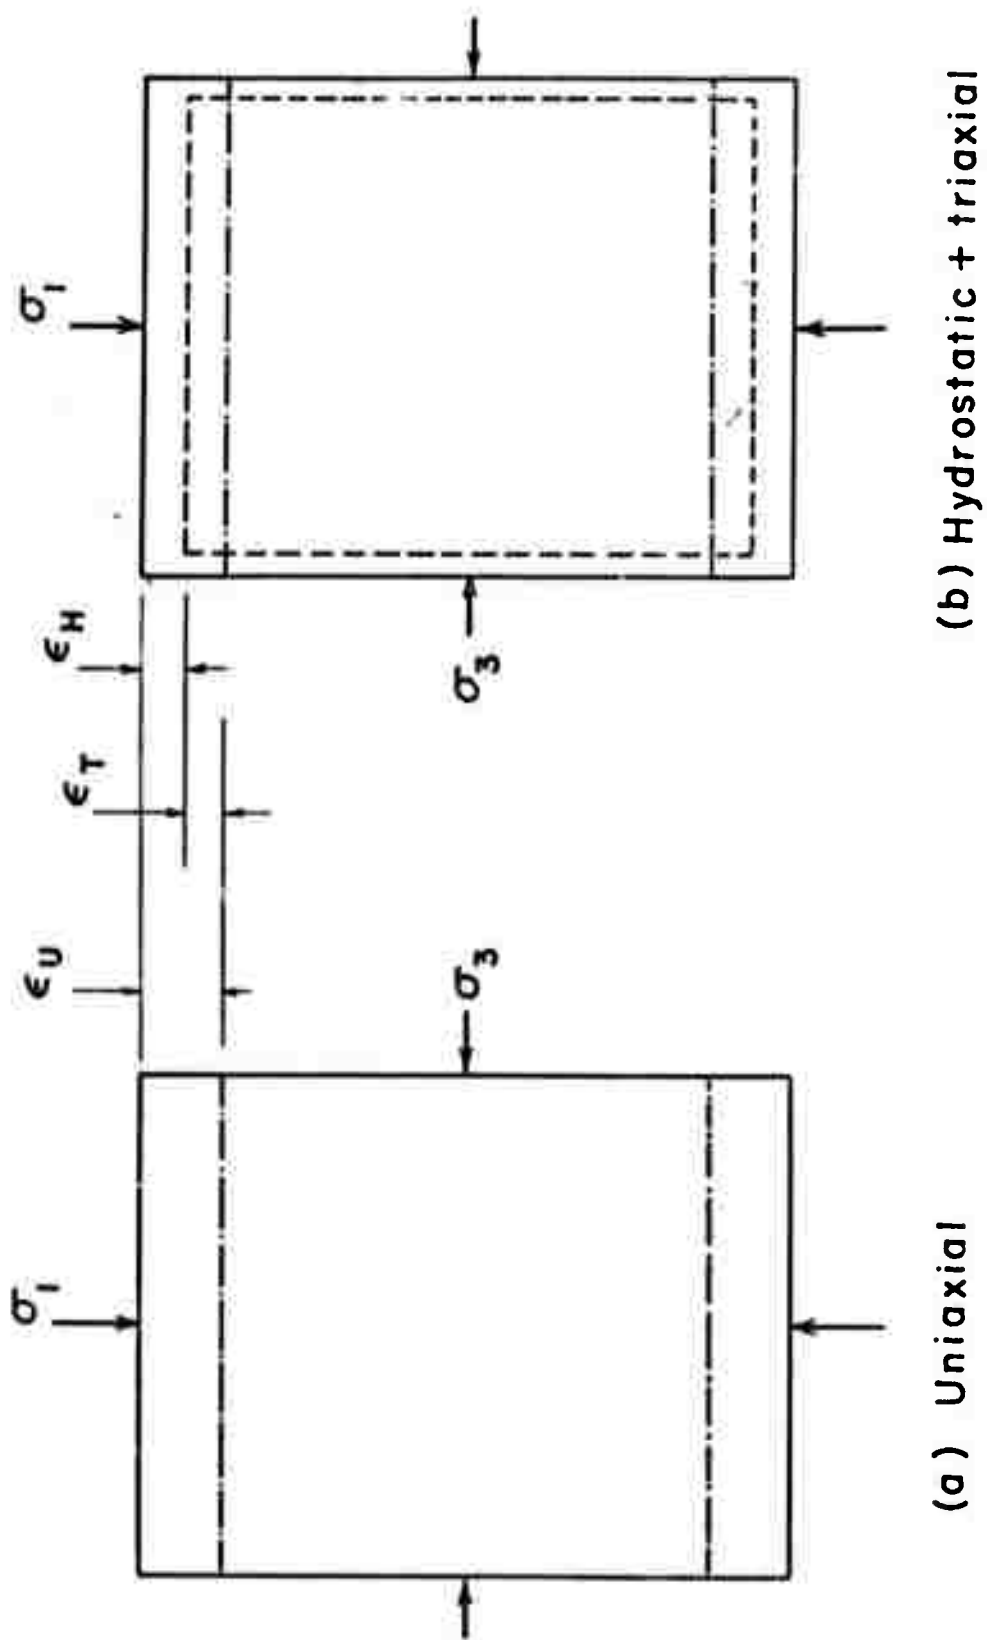
Figure 17. Sample Details for Uniaxial-Strain Experiment.

37
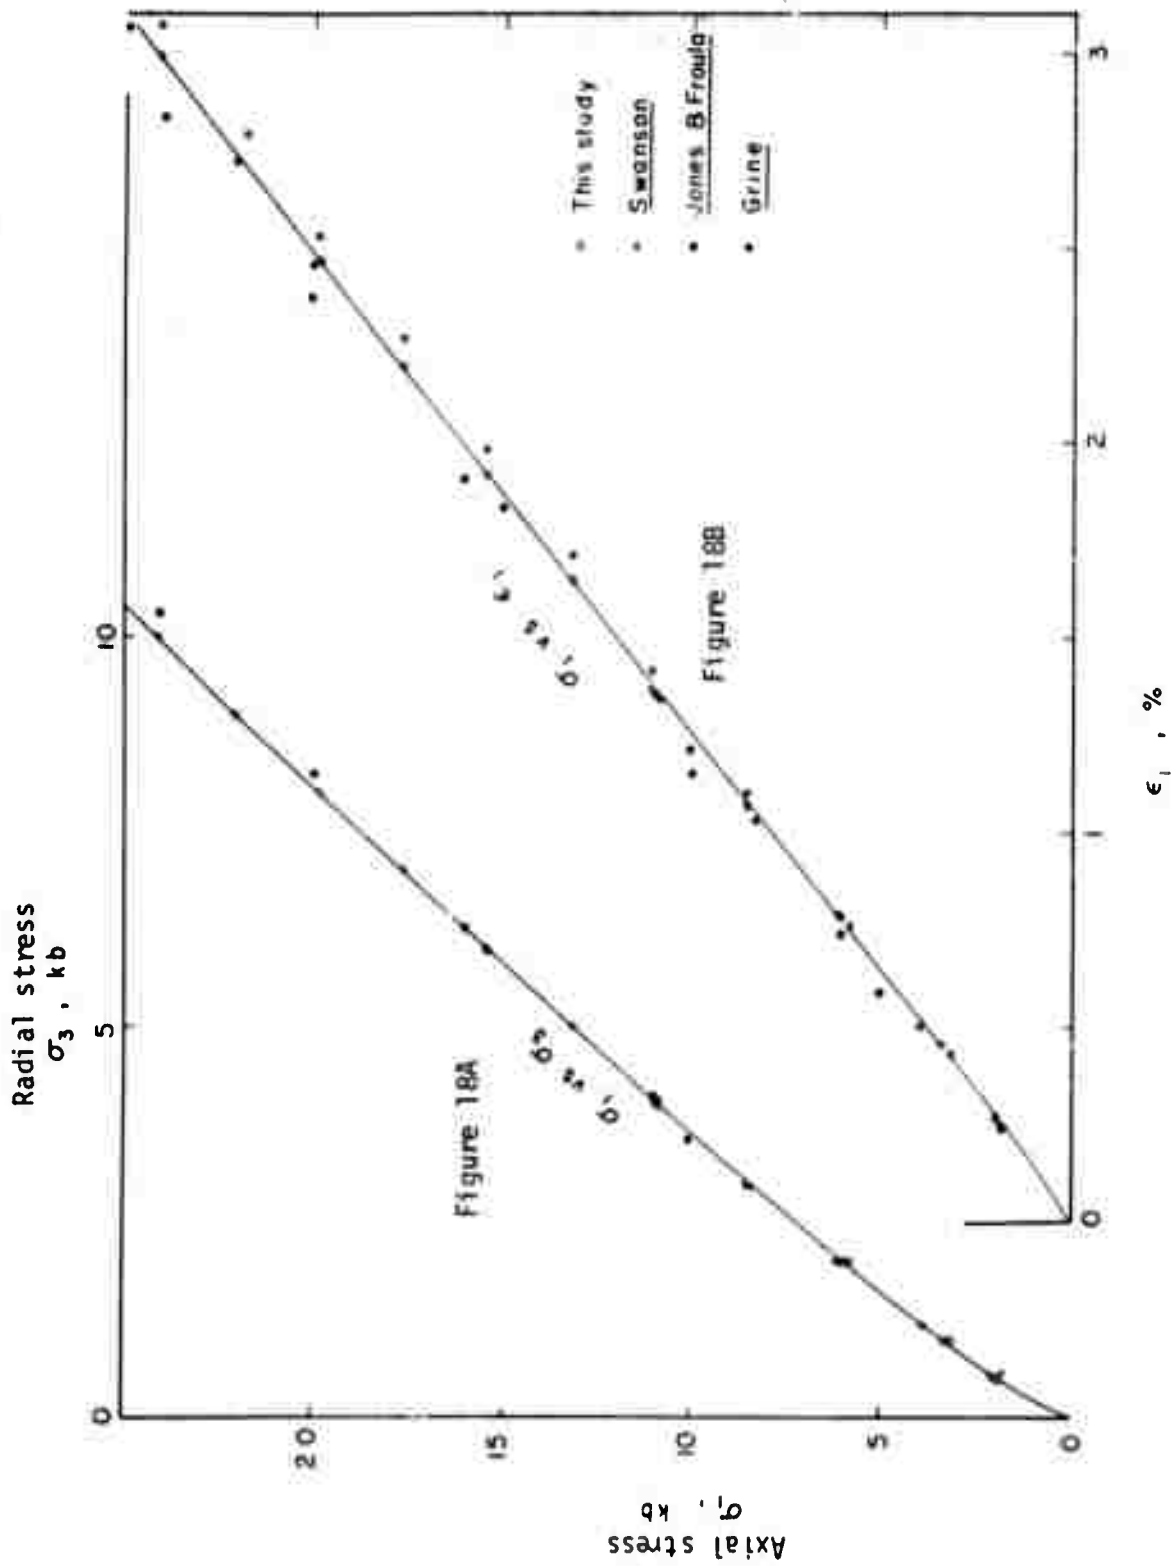
Figure 18. A: Radial and Axial Stresses for Westerly Granite.
B: Comparison of Uniaxial Strain and Shock Data.

We have compared ours with data given by Jones and Froula, Figure 18b, and there is very close agreement beyond a few minor differences near the beginning of the curve even though the strain rates in these two experiments are enormously different. The strain rates in our experiments are in the order of 10^{-4}; those in the Jones and Froula experiment are probably in excess of 10^3. So, essentially, the stress-strain behavior is the same in these two experiments.

I should add that comparison of work done at Utah and work done at MIT on the same rock so far shows very good agreement.

The final point is that for the second class of rocks, those materials in which there is porosity, there is inelastic, nonrecoverable deformation. For example, in Solenhofen limestone above about 6 kb you start getting a permanent compaction of the material. There we could see significant differences between shock and static uniaxial results, namely, that in the static, the slower experiment, there is more compaction, a greater volume decrease at a given stress than in the high strain rate shock experiment.

CHAIRMAN SIMMONS: Thanks, Bill.

MR. BROWN: Wayne Brown, University of Utah.

The curve that Brace just showed is very similar to our curve, and the significant thing is that the fracture curve is at best parallel. It seems to indicate divergence. The point is that we have not broken a rock in uniaxial strain. I think that is also the situation Brace found. The fracture curve which was shown seems to be somewhat universal and independent of the loading paths. We have tried different loading paths but constant confining pressure and also find that the curve seems to be universal.

MR. CHERRY: Wouldn't it be easy to reduce the material strength by saturating it?

MR. BROWN: The path has to change direction; it may very well do that.

MR. BRACE: These stresses are effective stresses that we show in a plot like this so if there were fluid pressure then we would have already subtracted it. Our general experience has been such that when we use effective pressure this completely skips the effects of the fluids on both of these curves within certain limits.

MR. HANDIN: I think the point is though that you achieve maximum capability of the strength of your apparatus. By saturating the rock, you can get closer to the fracture line.

MR. CHERRY: That is right. Reduce the strength or make the rock weaker. Why test the strongest rock you have in uniaxial compression when you know you are not going to fail?

39

CHAIRMAN SIMMONS: Are there other comments from the audience on this particular phase of the discussion?

MR. MCFARLAND: I want to make one comment. It seems that most of these experiments are stress controlled, and the problems of measuring the stresses are not going to come out of these in tests where you failed. For known stress you have a strain, and that is an indeterminate problem.

Perhaps if you can do controlled strain experiments to follow the stress you might get some more direct information.

MR. CHERRY: Yes, I think that is what I said.

MR. MCFARLAND: Well, the point is that these are stress-controlled experiments that we run.

MR. CHERRY: Yes, I know.

MR. GODFREY: I think the problem is--I don't know how to solve it--that when the rock begins to fail it is confined in a different way when a fluid confines it and there are resistances that build up that are not constant later. As the thing starts to fail along a fracture plane, the stresses, the lateral stresses, are not equal and they are preventing it from moving catastrophically. Although I don't know how to solve this problem, these are the kind of unloading data you would like to get.

MR. TRULIO: I think the immediate practical value of a uniaxial strain test lies in the fact that for tamped explosions, hopefully, you have a symmetric field, and that shock is uniaxial. It does subject the material to a uniaxial strain and at least you understand what happens to the material on the first shock.

CHAIRMAN SIMMONS: You mean it is radial, don't you; it is one dimensional.

MR. TRULIO: Semisymmetric.

CHAIRMAN SIMMONS: Is that equivalent to uniaxial?

MR. TRULIO: Yes, at the shock front. Behind that, though, of course there are specific divergence effects and those are very important. But you want to understand each part of the process for which we have to predict motion. I think one of the things that would be important, too, is to know if you did take the material to uniaxial failure, would it crack or would it powder? That information might have a significant bearing on whether gases have an important influence on the way pressure varies as a function of time in a cavity after an explosion. If you have radial cracks that is one thing; material powder is another.

MR. GREEN: Sid Green, General Motors.

40

I think it is worth commenting that we do have some profiles in addition to the classical profiles. We saw a break in the stress-time profile at about 13 to 15 kb which we had interpreted as yielding, but we are not so sure, and partly due to the insistence of Jones this may not be yield at all. We also have a shock that goes above 30 kb. All these are in Westerly granite.

MR. MADDEN: I have often made electrical measurements on samples, and there is obviously a very strong bias in this where the field measurements are almost always more conductive than the laboratory measurements.

The question referred to somewhat in Handin's talk is to what extent this bias enters into the mechanical problems that are of interest to this conference? Clearly in the electrical experiments I think the explanation would be that you never have in the laboratory samples from an adequate sampling of the cracks in place because certain cracks are left behind and never sampled in a small sample.

The question is to what extent is a similar bias important in the mechanical problems when we use the laboratory data to understand the actual field behavior?

CHAIRMAN SIMMONS: Do you want to answer that question, John? You might brief us on how well the lab data matches the field data.

MR. HANDIN: Obviously not very well, that is the problem. At this stage, I think, we are attempting, as best we can with limited sample size, to simulate the real world and particularly the defects to which I think you are referring. Then we hope that when we make the extrapolation of the larger scale that we are not overlooking something serious.

The only alternative to this procedure that I know is an in situ test which, in the case of strength properties, is very difficult. And, in fact, an in situ strength test at great depths where a rock is under high confining pressure has not yet been achieved. This is a disadvantage compared to your business. While your field tests may be more expensive, they are not in principle very much more difficult to make than the lab measurements. Ours are; because they are, all we can do in the laboratory is try to simulate the natural defects as best we can with full recognition that we are off in scale by orders of magnitude.

MR. BRACE: There is some limited work on large scale uniaxial loading, particularly in South Africa. They have worked with blocks of coal up to, I think, ten feet on a side and then they have the possibility of comparing field and laboratory at least in uniaxial loading.

Also, hard rock and norite, I think, were done there. The upper limit of size was about three feet on the side. The general results of these tests show that, as you would expect, the elastic properties are normally affected by the joints and are lower relative

to the values of the laboratory sized samples. The extensions aren't too different, particularly for hard rock. I believe about ten to twenty percent.

MR. SNOW: David Snow, Colorado School of Mines.

I wish to suggest two possible ways to get a larger scale effect. First, in the field you can induce uniaxial stress changes by lowering the water table. For instance, one can draw down the water uniformly. Even if you have plane-strain situation, you can measure the horizontal stress by some instrument.

Another approach that I propose is that one measure the mechanical properties, that is the stiffness, both normal and tangential, on natural materials tested in the laboratory, that is, naturally fractured cores with strain measured across these cracks as well as the intact portions of these cores. This would give you the necessary mechanical attributes of, let's say, a two-dimensional fractured medium of any size that you want to examine. Then, of course, you impose the boundary of no lateral strain and compute the change in the lateral stress that would correspond to a change in uniaxial stress load.

CHAIRMAN SIMMONS: Jack Healy has a short discussion to present.

MR. HEALY: Barry Raleigh asked me to come here in his absence and present some of the comments that he would like to have made, and I have some of my own. I think they are best handled in a short presentation.

Basically we have two things to say. One tends to simplify the problem, I think, and the other tends to complicate it. First of all, we have attempted to match near-source measurements from particularly the Sampson-Sterling nuclear experiments with seismic data.

The primary part of the problem that one would be interested in from a seismic or teleseismic point of view is that portion near the shot where the stress waves become elastic. It appears that we can model this physically by a very simple model, namely, a step function in pressure inside a spherical cavity. By superimposing this model on the data, we can model both the near-field displacements and spectral ratios for seismic waves recorded at a distance.

Now, this is a bit surprising, but I am going to show you some data that tends to justify this model. If this model is as good as it appears to be, it means that from the point of view of decoupling or the seismic observations alone, the details of what goes on inside of the cavity may not be very important, except insofar as they contribute to storing energy in heat rather than in momentum during the physical process that takes place.

The other comment that we have is one which stems directly from John Handin's remarks this morning and Raleigh and Curby's

42

measurements at the test site. If these statements about rock strength are correct, and if our observations about the state of stress at the Nevada Test Site are correct, then there is no possibility that we can see of evading major hydrofracturing following a nuclear explosion. This makes it absolutely necessary to do a three-dimensional calculation to study such problems as containment and also to understand the long-period radiation from these explosions.

[The slides used in support of Dr. Healy's statements were not available for inclusion in these proceedings.] In the first slide, we see the positions from Sampson-Sterling. Sampson-Sterling was about 2,000-ft deep, as I recall, and there were gages placed on radii both above and below, and particularly laterally with, the shot. Though we needn't concentrate on it too much, these are indicating velocity gages to give us ground displacements near the shot and I show quite a few of these. The point I want to make is that we can match these shapes of displacement curves with this very simple physical model which predicts the smooth line going through the data. We have allowed the apparent elastic radius to vary in matching, but you will notice that they are not inconsistent. There is 156 on the top, then it goes to 143, 137, 143. The principal point here is that this simple model matches the data better than the experimental accuracy. In other words, consistency of the data with the model is better than the internal consistency of the data.

The next slide shows some more of the same data, and I want to call your attention particularly to the two bottom traces, which are at 659 m and 744 m. They both give an apparent elastic radius of 174 m. These were the farthest measurements from the shock and are above the point at which we believe the propagation finally turned from plastic to elastic. Again, look at the remarkable match with this surprisingly simple model.

The only substantial deviation that isn't explained at all physically is the very tiny forerunner which is one of the well known elastic precursors. The match for Sterling was not as good, but again the two farthest measurements are fairly well matched, and the deviations, the big negative displacements and the two closest ones are not predicted by any model. They might possibly be errors in the gage measurements or they might be related to some kind of cavity collapse following Sterling.

Anyway, my main point is that I think these measurements show there are some very simple approximate calculations that might be done in the decoupling problem. In the next slide we have gone ahead to use this concept to interpret the seismic failure. These are Sampson and Sterling seismic spectra, smoothed and unsmoothed. I only want to show you the difference in character, Sampson on the left and Sterling on the right. The top is unsmoothed and the right is smoothed.

43

The next slide shows the spectral ratios of these two for the number of different sites which were recorded. Once we accept this simple model, first proposed by Shulips some 30 years ago, we can then calculate from the seismic data the apparent elastic radius. We calculate an apparent elastic radius for Sampson of 169 m which compares with 174 m from the minimum, and one for Sterling of 29, which compares with 130-140 m. Sterling isn't quite as stable. The pressure ratio is about 4.15, which compared to about 5.

If we use this simple model it leads us to our next conclusion. This is predominantly an outward motion around this explosion. The details of how you get into that step control the bending in this spectrum. They probably are rather limited by simple geometric conditions, i.e., the radius at which this becomes elastic is a powerful control on the nature of this displacement motion, and it is to some extent independent of the physical behavior of the rock materials.

We calculate for Sampson at that elastic radius, a pressure of about 400 bars, and this requires a much higher pressure within that cavity. The natural cavity was 17 m, which presumably was 20 m at the time this stress wave passed by. So the pressure inside that cavity had to be much greater than 400 bars. Now what happens in a cavity at 2,000 ft or so where, I think, the overburden pressure was about 80 bars with some half a kilobar or so of pressure in this? All that we have heard this morning says that there must be hydrofracture occurring.

In salt the hydrofracture is probably going to be vertical if there is any horizontal stress at all. We don't have direct evidence for this in Mississippi, but we do have quite a bit of evidence at the Nevada Test Site.

Barry Raleigh attempted to measure stress inside a hole that was later used for a nuclear explosion. The chambers were big enough to descend in an elevator and work in the base of them with drilling equipment. They used an old technique of drilling several holes into the wall rock, mounting or pasting a transducer on the surface of the rock at the base of the hole, and then in steps making measurements of the strains as they have penetrated the wall.

This particular method was tested by Earl Hopkins and calibrated in the laboratory. They obtained the following results: The σ_{zz}, or the vertical stress, was about 4,000 psi; the σ_{xx} for one of the horizontal stresses was about 1500; and the σ_{yy} was about 2,000 psi. In other words, the horizontal compressive stresses were substantially less than the vertical, which would be the case for normal faulting.

Unfortunately, they had a drilling failure and they didn't get enough data, so they don't believe they have accurately determined the principal stress. But they do have this value of less than 2,000 psi, which is about 70 bars, and that is in good agreement with results in seismic measurements.

44

What does this imply? It implies that following the detonation of this shot there should be a major vertical hydrofracture through this region, and it also implies that the rock around the shot should be accompanied by normal faulting.

I have some slides to support these contentions. Next slide, please. This is a strain meter, a vertical strain meter, installed in a 4,000-ft deep hole. These were operated both for Jorum and Handley, and both were very similar. The vertical displacement is on the bottom scale. Notice that at the time of the shot there was a shock; there was clearly a major explosion. Now for the Jorum shot, Cal Tech, and I believe also the Colorado School of Mines, has shown a delay here that took place in approximately 30 min--a very long delay. Barry and others were looking for this. But both for Jorum and Handley following this major pulse there is a delay here that lasts only about 70 sec. It is impossible to get this kind of delay from any thermal cooling in the cavity. It is also inconsistent with any of the simple seismic or plastic models around the source. One good explanation for why this might occur is the major hydrofracture which relieves the cavity pressure and allows it to decrease.

Let me run through the next slide. These are earthquake solutions associated with Benham. Notice there is one north-south, but over near the source there is predominantly dip-slip motion which is again in agreement with this in situ stress measurement. These are the tension and compression axes for these fault plane solutions. The open circles are the tension axis indicating a tension northwest and southwest in direction, which is again consistent with this other measurement.

In summary, I think these two points should be considered: One is the possibility of understanding the first compression wave may not be as complicated, because we can summarize some of the complexities within the cavity and neglect them, looking only at longer term displacement along the elastic radii. The second seems to me a necessity of looking carefully into this matter of whether or not we are hydrofracturing the zones around the nuclear explosion.

MR. GODFREY: In line with John's comments about a cheap in situ test, I might just point out that there was a program we were involved in where the specifications required measuring a dynamic, in situ modulus at a specified loading state. We didn't know how to do it, and we devised an experiment that involved burying a sphere of explosives, about a ton, which is about as small as possible to still get data that is significant with respect to the jointing, spacing, and other things. You put instrumentation at various radii and you measure the loading and unloading. We were proposing to do 200 or 300 of these tests, but that part of the program didn't ever mature. If John thinks $100,000 is cheap, that is about the cost of one experiment.

CHAIRMAN SIMMONS: Jack, what is the status of hydrofracturing in hard rocks, like granites? Has anyone ever tried this? I am familiar with the work of the oil industry on sedimentary rocks.

MR. HEALY: We think that is what happened in Denver, but there is really not, to my knowledge, a great deal of engineering experience with hydrofracing and the accompanying measures down hole to explore these measurements. But it is the only way that you can explain the amount of fluid that was pumped in the Denver well, for example, or the rate of change of permeability. Another point is we don't understand, but we are suspicious of, the very low fluid pressures in the crystalline rock beneath the well. It would seem to me that there is a suggestion that hydrofracing can play a major role in fluid propagation in crystalline rock, even under natural conditions in some cases.

For example, at Denver if you ran the Platte River over to the hole, which wouldn't be any trick at all, it would run down that hole as nearly as we can tell.

CHAIRMAN SIMMONS: I have forgotten how big the river is there.

MR. HEALY: It is pretty big. It could be very exciting.

MR. PRATT: I would like to mention that we are running a series of in situ tests at Cedar City, which we hope are cheap, and we are using a configuration of cutting out a triangular shaped prism, cutting slots here and here and one at the end. Then we load it with either a series of flat jacks or explosively if we can do it, dynamically instrumenting the top and side. We were in the field experimenting last week.

We hope by early fall to have some information. Because the tests are 8 in. on the side we can excavate the rock and take it into the lab and crush it in a large testing facility. We are trying to make the transition from the lab to the field and trying to include parameters that are inherent in all the rocks. So we are hopeful, and it is cheap.

MR. ATCHISON: Tom Atchison, Bureau of Mines, Minneapolis.

I would like to comment on Jack's remark of the simple, cheap solution. In the work the Bureau has done in smaller scale studies, both crater formation and stress pulse work, including wholly coupled and decoupled charges, we found that we could not only get pressure curves or strain-pulse curves with amplitude conforming to the theoretical step-function predictions when we use the elastic radius as our source, but we also could confirm the frequencies.

Of course the trick is determining the elastic radius, and we actually worked backwards in the beginning, starting from the frequencies that we have developed in our strain and pulse measurements for various decouplings. From this we were able to get a correlation between

46

MR. HEALY: There wasn't in Sampson; there was a very close match. There are questions about whether an individual measurement is good to within a factor of two in the near field. There is much scatter within the data. I am saying that our model matches the observations as well as any other model, and because it is so simple I prefer to use it until we find another model that matches better.

MR. TRULIO: By a model, you mean merely that you assume a step function of the elastic radius and from there it is just a matter of elastic propagation of that particular pulse, a displacement of that function.

MR. HEALY: We don't assume, we demonstrate that we can match the ground displacement by an elastic solution. It turns out that the elastic solution is a step function inside a specified cavity. That is what all those figures are that we showed.

MR. TRULIO: If we were to postulate, for example, a given yield in a given size cavity, how would you use the model to calculate the signal? What are the input and output?

MR. HEALY: Suppose we postulate a given yield in the salt belt. We could take that radius of 174 m and that will scale, then, as the cube root of the charge size.

MR. TRULIO: That certainly isn't true for a fixed cavity radius.

MR. HEALY: This is a tamped shot you are talking about.

MR. TRULIO: No, just an arbitrary configuration. Let's assume it is symmetric for the sake of argument.

MR. HEALY: If we could assume that the cavity was made without changing the strength of the salt, then we calculate from the model that these become elastic at approximately 400 bars, with the pressure wave going out, which may be a little different than the long-term behavior. The long-term behavior probably can't be elastic. But this very short pulse appeared to be elastic when it reached a level of about 400 bars. So we can calculate that when the pressure out of the cavity reaches a level of 400 bars it will be elastic. Then I need to know the equation of state of the cavity and what the cavity pressures will be. We can possibly extrapolate those pressures and determine the elastic region of the radius. From that we know everything about the spectrum of the source, with the exception that it is not perfect. We are arguing about a factor of two maybe, but not an order of magnitude.

MR. TRULIO: But you are not saying you can calculate an elastic radius for an arbitrary shot, are you?

MR. HEALY: It depends on the shot, which can be measured from another shot or laboratory measurement.

MR. TRULIO: The elastic radius also depends on the mechanical details.

MR. HEALY: Not in this light. It only depends on the equation of state of the salt at that pressure. It doesn't care what happened in that cavity. It only depends on what happens as the stress wave passes that point.

MR. TRULIO: I guess I am not clear again. If I understand you, you are saying that you can predict the radius from any given circle configuration--by predict I mean not do an experiment and measure it. From your model I am not clear how you do that.

MR. HEALY: You need some experimental data. But we are saying that you don't need the whole equation of state. You only need to have the piece of data that tells you when the waves become elastic.

MR. TRULIO: That is a pretty involved thing to calculate.

MR. HEALY: Then we will measure it. Put a small shot in the salt and measure it.

MR. TRULIO: I can see this is valuable for diagnostic purposes, provided that there may be limitations on it.

CHAIRMAN SIMMONS: You keep saying salt. Are your data only for the salt shot?

MR. HEALY: The data I presented comes from the Sampson Sterling.

CHAIRMAN SIMMONS: Are you willing to extrapolate your model, then, to hard rocks?

MR. HEALY: I think it is suggestive, and Jack Murphy can tell us more than I can about that.

MR. MURPHY: We have used something of this sort for rhyolite and tuffs at the test site, yes. In other words, we used initial measurements for the elastic radius and pressure and profile at that radius to compute the seismic forcing function. Given a measurement at a distant station, we would have scale from that and scale from the source function. This, of course, would vary from media to media, and the initial measurement would determine the elastic radius, and possibly the pressure profile might change from media to media too.

CHAIRMAN SIMMONS: As I understand this discussion, Jack says--and some others here seem to disagree--that given very simple models of the equation of state, namely, the pressure at which the material becomes elastic at the appropriate frequency, then you really don't need to worry about the details of strength; that you can calculate everything from your very simple model and it matches the data for the salt very well, is that correct?

MR. HEALY: That is pretty close. I am not saying anything terribly surprising or profound. The spectra of the seismic wave is controlled predominantly--the factor of two is in argument--by the elastic radius or the apparent elastic radius. This can be measured in a number of ways, either in situ by firing a small shot or by laboratory measurements, and this tells us about the spectra of that shot.

MR. PHINNEY: There is one point where we are a little confused. We are taking elastic radius to be any radius out to this point which you are describing as elastic radius. Any point in the elastic region can be used as a point on which you would specify initial conditions of source function and from which you could then in principle compute the response at any other point in the elastic region. Are you saying that you actually provide the elastic radius?

MR. HEALY: The elastic radius. This would give the best fit for step function in pressure.

MR. PHINNEY: So you have a one-parameter method?

MR. HEALY: So it is an apparent elastic radius, calculated to fit this simple model that might vary a little. Its physical meaning is not clear. It is, of course, a gradational radius, but the numbers that come out of the analysis are rather precise.

MR. PHINNEY: This gets into the matter of source function now. If I try to determine what happens in the real problem, and I take my far-field data and analytically continue it inward, then you are saying I reach a certain point where there is a step I might actually go. You have essentially computed the radius of which there is a step, but that is not necessarily ....

MR. HEALY: No, no. You can take your far-field data and analytically continue them, at which point it will begin to depart from the actual measured values because the model is no longer behaving elastically.

MR. PHINNEY: How do I know this?

MR. HEALY: Because I showed you the profiles demonstrating that this apparent elastic radius is increasing as you move out from this source. The pulse, in other words, is broadening; the primary plastic effect seems to be broadening this pulse or making the equivalent elastic radius appear larger, which is physically very satisfying to me. It is not perfect.

MR. PHINNEY: I guess that this apparent elastic radius is not clearly what these gentlemen want.

MR. ROTENBERG: The radius you are talking about is an input to the model, not an output from the model.

MR. HEALY: We input either the near-field displacement measurements or the spectral ratio of the Sampson-Sterling and we came out with the elastic radius.

MR. ROTENBERG: I think the confusion results from your calling it an elastic radius. It really isn't.

MR. HEALY: I call it an apparent elastic radius. I was very careful about that terminology. It does appear to be the elastic radius, I don't know how you know that it really isn't. Because at this radius you do see that the elastic precursor which we showed you that continues to develop, the velocity of this main pulse is lower than the elastic velocity. There is lots of evidence to suggest that in this ore case which I showed it was inelastic propagation out to that radius.

MR. PHINNEY: This may be different from what you have done, but it is not clear to me. Has anyone actually taken the data at a near-in point and used this data as the source function, as the initial conditions for continuation of prediction of data at a greater distance and assuming that we are in the linear region?

MR. HEALY: Ask Ted Cherry.

MR. CHERRY: I do try to predict the reduced displacement potential functions for all the shots that we do. I found that the reduced displacement potential is affected seriously by the properties of the material, such as the amount of hysteresis, strength, and so forth. We have tried this technique and had some success. You can vary the parameters in the code and you finally match the observed potential. For a model to be a real model it has to have some prediction capability. I think we quarrel here about the advisability or the cost of actually doing a shot to determine the seismic efficiency of the material; we have all felt that it is fairly expensive.

The thing that we are trying to do, it seems to me, is to devise a suitable testing procedure, pre-shot in the laboratory and maybe some corresponding in situ logging tests to adequately describe what a material is going to do and how it is going to behave so that at that point you can predict the seismic effects. It seems to me that was the issue we were addressing ourselves to here. Just taking the shot data and saying, "Well, we can match this with a step function and pressure at some radius," gives some insight into what has happened. But with this approach, maybe you start to worry about how good your predictions are going to be on the next shot in a different environment, and where the elastic radius is going to be and what sort of pressure you are going to have to apply to that elastic radius to predict the observed displacement at a point.

CHAIRMAN SIMMONS: I expect that some of this same discussion will be appropriate following the presentation on code calculations. I would like to turn to something else, with the panel's and the audience's approval.

MR. CHERRY: I would like to return to a point that Brown brought up a while ago. I believe he said that the loading curve in stress-strain space did not intersect the fracture curve, the fracture line for most material. If that is true, it seems to me that maybe we could conclude that fracture is not going to be important in wave propagation in general. The question is, do we use these fracture data in our codes, and to what extent are we dependent on that research?

MR. TRULIO: If that is true, in the first place it just applies to uniaxial loading and that can only be assumed to take place at a shock found in the kind of problems we are really interested in. Fracturing can occur at any stage later in the motion. Failure could occur under much more complicated conditions of stress, and those are observed to happen in fact in code calculation. That is why, as I think I tried to say before, this is just one piece of the problem, but a good one.

MR. CHERRY: I certainly have to agree with that in principle, but I am just wondering does it make any difference quantitatively?

MR. TRULIO: Yes, it does.

MR. GODFREY: I don't think one has to worry about what you suggested, because with the fracture curves that I have seen, if you go up high enough they become almost horizontal ultimately. Whereas the uniaxial loading curve, if you follow it up high enough, becomes convex upwards. They must intersect somewhere.

CHAIRMAN SIMMONS: Bill Brace, would you like to comment?

MR. BRACE: I think Wayne Brown has probably done more of these experimental studies of uniaxial strain than I have. Would you agree with that last statement? Will it have to intersect somewhere?

MR. BROWN: It could be. But we haven't seen this.

MR. GODFREY: I will retract that statement, because I was talking about something that does go up.

MR. BROWN: We have seen the curves become parallel and start to diverge. Although we think maybe there is some kind of intersection if you get far enough.

MR. VANNING: It might be very well that the cracks are already there anyway. We do have large rocks. Maybe that is the reason we don't have this problem. The cracks are there and all they do is allow the material to slide.

MR. STEPHENS: On shock loading you have a uniaxial experiment and a need to get a limit on every rock eventually. This implies a one-dimensional strain experiment has achieved failure. I know of no obvious exceptions to this rule.

52

MR. MOORE: Henry Moore, Geological Survey.

I wonder what sort of geologic observations were going along with these programs. My experience has been with missile-impact craters. Immediately beneath the floor of the crater we find mixed breccia composed of projectile and sheared and compressed target material. As you move outward from the impact point we run into a region where the material appears to have flowed. Beyond that we run into a set of fractures and then this dies away.

In addition, you run into scaling problems for the missile-impact craters which are very small and might be only about 18 ft or so across. The spacings for the size of the fractured pieces are a matter of inches across, whereas if you go to a place like Teapot Test Site the spacing between the fractures is quite large.

MR. MC FARLAND: When you reach the stress at which brittle fracture would occur, you very likely do not have a shock front anyway. That is a point to consider. Another thing is that the uniaxial strengths have indicated in some instances they do intersect the triaxial failure envelope and in other instances they do not. This is the sort of case that Ted was talking about earlier. In that instance apparently it has some material test.

CHAIRMAN SIMMONS: I have a distinct impression that the study of the property of rocks ranges over a whole spectrum. At one end is the treatment of rocks as a very simple-minded, isotropic, single-phase, very well behaved material, which is probably the approach of people who know nothing about rocks; and at the other end is the very detailed approach of the petrographer who spends a lifetime studying one thin section. It seems to me somewhere in between we ought to have a careful look at tying these two things together. I have a suspicion that one of the best places to do this is in the examination of microstructures and the effect not only of the elastic properties but the inelastic properties. There are a very few people who have done anything along these lines. I think the people in ceramics have found there is much information in the microfracture of their materials that they can relate to strength and can relate to elastic property.

So I might ask a question: Do you know of anyone who is doing this kind of thing for us in rock mechanics? I don't.

MR. GRINE: We are doing several projects on which we are studying the effects of both microfractures and hydrofractures on uniaxial strain dynamic experiments. We have microfractured Westerly granite, for instance, by subjecting it statically to unconfined tests. Then we have done uniaxial strain with a gas gun on this rock as well as on rock taken from adjacent locations. On the dynamic loading curves we were not able to see any significant difference in the whole shape of the loading, including the elastic limit from a microfracture. That is certainly part of the rock texture.

53

We are doing the same things on rocks that have multiple components, like granite as compared to quartzite. We see quite a difference in the shape of the loading curves in whether or not one obtains a steady state.

In general, in the single component rocks like quartzite, we have studied specimens up to 4-in. thick. These experiments are not repeated. In granite we do not see a steady state attained. The rise time continues to increase over the thickness that we so far studied, which is only up to an inch, but we are going up to a foot in the near future.

MR. HANDIN: I am not sure I understand your question. But I don't think there is any rock testing man in the audience who doesn't do his work in conjunction with very careful microscopic studies of deformation mechanism. Perhaps that is not what you had in mind.

CHAIRMAN SIMMONS: Specifically I had in mind from my own work, we think we can explain the behavior of the compressional velocity in terms of as yet perhaps unidentified microcracks which are much smaller than the ones this gentleman is talking about. In addition, there are always elastic mismatches at grain boundaries, and it is this kind of very small microscopic detail that I was raising a question as to whether anyone really looks at and tries to understand in detail, rather than what you are suggesting. There is no question that everyone is careful to look at the sections of his rock before and after deformation. I didn't mean to impugn your reliability as an observer. I was trying to raise a question. I think what happens between grains elastically is a very complicated relationship. I don't even know how to set the boundary conditions myself.

MR. HANDIN: I am not claiming we understand all these things, but I do claim we are trying.

CHAIRMAN SIMMONS: That is all I am asking.

MR. KNAUSS: I heard a question from Dr. Cherry essentially asking what the experimentalist's evidence has to say for the input to his computer program. It occurred to me that the experimental data presented today deal essentially with fracturing in simple laboratory tests where the fracturing occurs along fairly well defined faults and that the subsequent strain behavior incorporates not only the motion along the fault but also the deformation of the total specimen.

It seems to me when you want to do computer calculations you again deal with large rock masses, but these large rock masses contain several faults and the motion occurs mainly along faults.

• The question therefore arises, in a given mass of rock how many faults do you have? The motion that occurs along faults is the primary mechanism you may have to incorporate. I have not heard anyone

54

say anything about this. When you do this, although I have not, you end up with computer calculations possibly as gages to nonunique solutions, because then you begin to get into the area of, roughly speaking, elasticity.

MR. MOORE: In the crater we find sheared and compressed material. The original material might be 0.5 g/cm^3 and then compressed to about 2.2. If we look at the craters in water-saturated lakebed, we find no evidence for shearing and compression on the target area.

These kind of observations are pretty important. And this is accompanied by the fact that the craters in the water-saturated material are about six times larger by volume than their dry counterpart.

MR. BARRON: I want to get back for a moment to the question of models and point out where we stand now as to where we stood two years ago.

At present, we are certainly, from laboratory tests, beginning to get from such people as Wayne Brown at Utah and others pretty complete sets of data, admittedly on relatively small samples, but fairly complete, such as uniaxial tests, triaxial tests, proportional-loading tests, and pressure-volume tests.

What has evolved among various people is a capability of developing what you might call mathematical models for the large codes which, at least before you put them in the code, fit the data that come from the various tests in an overall picture. This is quite different from several years ago where you may have had one test and one measurement from a test and sort of threw a coin to figure out what the other parameters were going to be. At the same time, it seems to me, whereas this apparent elastic-radius approach is perhaps a good empirical approach to getting certain answers, if you want to apply an approach to different materials under different conditions, etc, number one, you have to start off with a mathematical model to do your calculations which will satisfy the experimental data, and, secondly, when you run it in a code you will have some hope of uniqueness.

You can by playing around with two or three of these parameters, in some cases get from one point to another point in the medium. This does not mean that you are solving the problem. You have to be very careful if you are going to use this approach from a general viewpoint. Therefore, my feeling is that one has to progress along the point where starting at the beginning one can calculate the fields at all distances. If you can, then you think you are not afraid to use it as your next problem.

MR. HEALY: Let me try this question again. The step source in a specific cavity has one unique thing about it: it is the simplest conceivable physical model that might match the data. It is just the simplest one.

55

We took that model and tried it against real data. We matched the Sampson-Sterling data as well as anyone has ever matched any data at the elastic region. Furthermore, the model is adequate to tell us what this means in terms of teleseismic decoupling.

Now, if you believe the work of Love and Lamb and others and the uniqueness they prove for the solution of the elastic wave equation, if you can fit the boundary conditions and match the displacement shape, then that is the only elastic solution to the problem. There is no other. It is not a question of having 40 different models.

We admit that there is a factor of two; we are not saying there is anything wrong with attacking this problem with large computers. It is certainly necessary. But I don't think we should resist the fact that a very simple idea gives us a lot of answers.

MR. BRACE: I wonder if I could come back to the question raised earlier about the behavior of the material with multiple faults and fractures. I was talking to Steve Swanson of Utah this morning. I think there was some very new work which he might mention just briefly.

MR. SWANSON: Yes. We compared some experiments that Byerlee and Brace ran at MIT with some we did at Utah in an entirely different experimental arrangement. We had a fragmented specimen, multiple-fault arrangement, and they had cut single faults carefully in specimens and calculated friction. Our data compare exactly. We can look at the maximum stress that this fragmented material can carry and it can be predicted right on by the measurements that they made with a completely different method. Also the deformation is reasonably similar.

MR. BRACE: Have you done uniaxial loading in this type of experiment?

MR. SWANSON: No, we haven't yet.

ON THE APPLICATION OF FINITE-DIFFERENCE METHODS
TO STUDY WAVE PROPAGATION
IN GEOLOGIC MATERIALS*

Henry F. Cooper, Jr.
Civil Engineering Division
Air Force Weapons Laboratory

Abstract

A review is presented of the capabilities of state-of-the-art finite-difference methods to understand late-time ground shock phenomena associated with explosions in various earth media. The key problem areas involve proper evaluation of numerical errors and valid determination of in situ material properties. Although examples of numerical errors suggest that one must apply code results with great care, the theoretical techniques can provide phenomenological understanding which, when coupled with meaningful in situ material property tests, can provide high-confidence predictions of explosively produced ground motions.

Introduction

My original charter for preparing this paper was to review state-of-the-art finite-difference procedures from the viewpoint of assessing code capabilities. Following several discussions with Dr. Black, Dr. Ruby, and Col. Russell of ARPA, my objective has expanded from presenting primarily a technical survey or review paper on computer codes to presenting ideas to stimulate discussion that hopefully will lead to an assessment of the technical community's capability to solve the detection problem in a finite and foreseeable time period. I am sure I can stimulate considerable discussion and heated debate of the virtues (or lack thereof) of the current theoretical procedures and what can be done to make these procedures "better." However, it is a more formidable task to provide the baseline discussion for positively addressing ARPA's specific and practical problem.

Perhaps the best place to start is with a statement of the problem. My understanding of the key objective is to: *Determine the late-time displacement produced by the detonation of underground explosions in all possible media for all possible source configurations.* The highest confidence solution is, of course, to test all possible configurations in all possible media and to obtain a direct empirical solution. This solution is completely impractical from the standpoint

---

*This paper was previously issued as technical report WLC-TR-70-171.

of time and funding restrictions, although a significant amount of data already exists. Thus, the real question is whether or not a practical solution of verified confidence can be found by a systematic interaction between theory and experiment.

It should be noted that the prediction of peak particle displacement is a considerably more formidable task than predicting peak particle velocities, stresses, and strains. Generally, these latter variables are controlled by the immediate response of the earth media to the shock wave passage, whereas peak displacements are governed by late-time behavior of the in situ rock. This late-time phenomena is controlled by reflections from even rather distant inhomogeneities and by the in situ behavior of the earth media.

In particular, the determination of in situ material properties is quite a key issue. It is difficult to see at present how one can, with confidence, make direct use of laboratory-determined material properties in estimating the late-time response of an in situ rock mass to an intense shock wave. One has to stretch his imagination in biasing, toward one extreme, the laboratory properties determined for granite to even come close to reproducing the displacements measured on contained bursts in granite (1-3)*. This point may be further demonstrated in another way by comparing data from underground experiments in Ranier Mesa tuff (a "soft" rock) and granite (a "hard" rock) as shown in Figure 19. Note that peak particle velocities differ by a factor of about 3.5, which is roughly consistent with the difference in seismic velocity (or peak wave speed, for that matter). This is also consistent with about an order of magnitude difference in confined modulus (ρc^2), which in turn is roughly what one would expect based on laboratory tests of small samples of granite and tuff. Note, however, that no great difference between the tuff and granite is indicated for peak particle displacements--a fact that would be surprising if one expected rock properties based on laboratory tests of small samples.

It should also be noted that the cavity displacements in French underground tests in granite were significantly smaller than those in experiments in NTS granite (4). Some have attributed this difference to the difference in water content--the French tests being in a much dryer rock. In any case, there is the suggestion that granite in one part of the world may behave somewhat differently than granite in another part of the world and further that displacements in some granite and some tuff are quite comparable.

In understanding close-in, late-time motions in rock, a high priority must be placed on the understanding of the motion along jointing surfaces. Figure 20 shows the permanent relative displacement along jointing planes between blocks of rock resulting from a close-by, high-explosive detonation.

---

*References are listed numerically on pages 110-112.

58
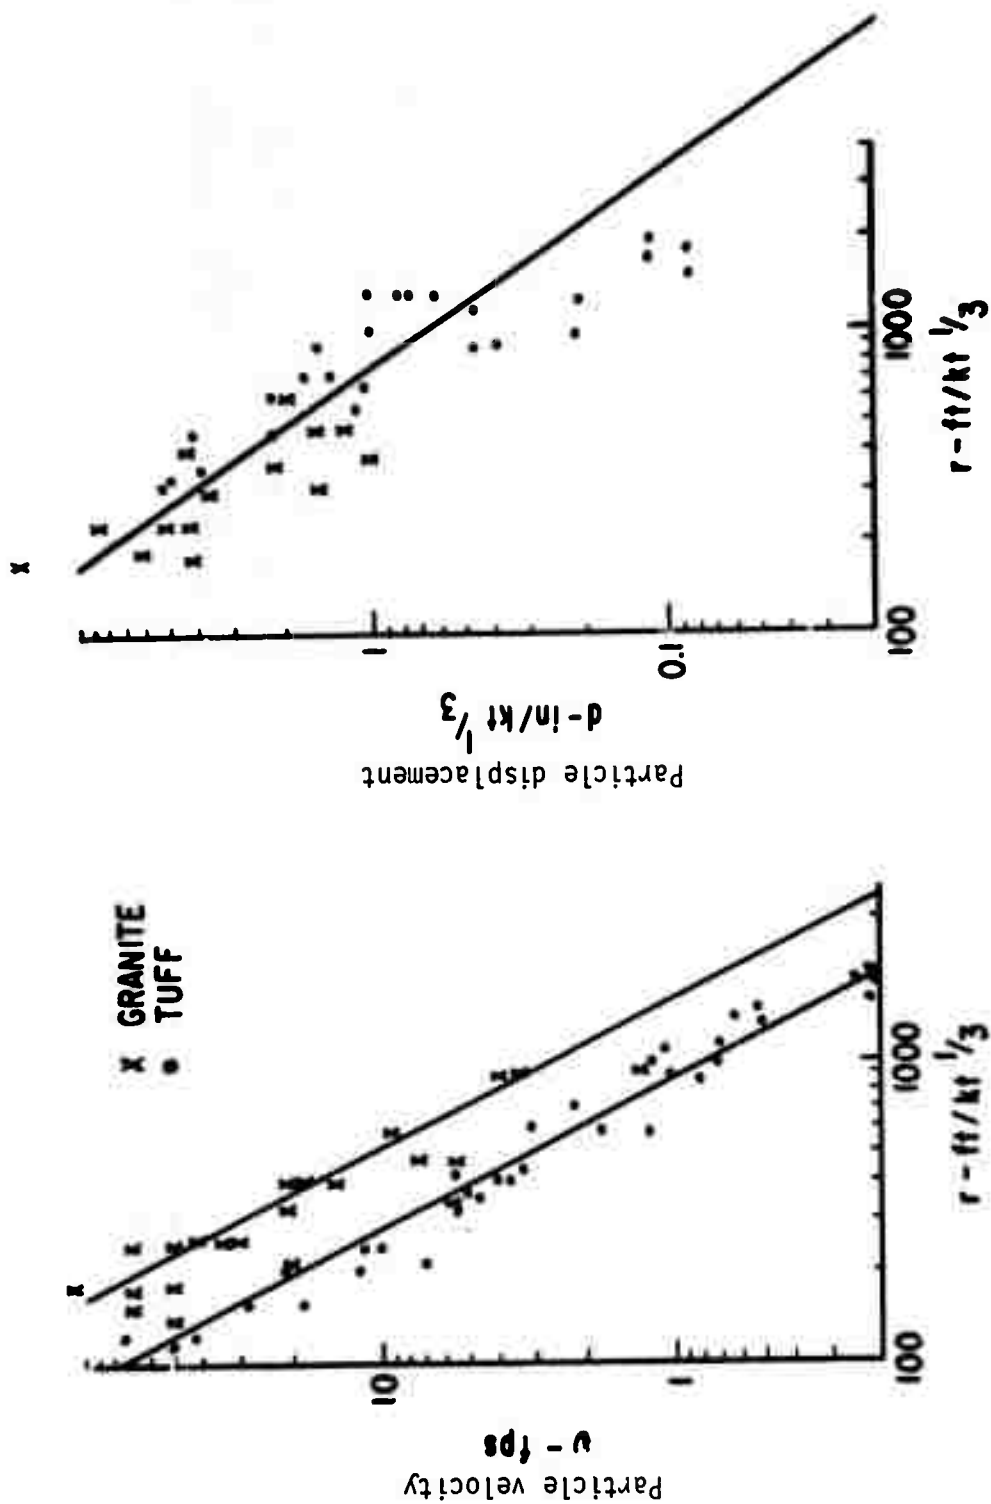
Figure 19. Ground Motions from Contained Bursts in Rock.

59

Figure 20.   Relative Motion Between Blocks of Rock Caused by Intense Ground Shock.

This hole was once a right-circular, smooth-walled cylinder. It seems intuitively clear that this observed phenomena will not result from calculations for homogeneous media described by material properties based on laboratory data. It also seems clear that the material filling the joints, how well the joints are healed, and other details not provided by laboratory material property tests may be tremendously important in influencing late-time phenomena.

Based on the above comments, it is reasonable to question whether or not we should expect to ever be able to calculate such late-time behavior from first principles. Second, it is also reasonable to question whether or not one can develop any credible procedure, based on both theory and experiments, that can satisfy ARPA's requirements. A credible procedure for predicting the results from a given test event either must include in its derivation data from similar previous events (in which case the prediction procedure can be purely empirical), or the procedure itself must be validated by tests sufficiently rigorous to establish high confidence in the extrapolation from a known data base to a new situation. In my opinion such confidence can only be achieved by consistent success in providing accurate pretest predictions of highly controlled experiments not included within the pretest data base.

60

While I believe that post-test calculations can significantly increase our understanding of basic phenomena, I do not consider them to be valid as building confidence in our ability to predict new and different situations.

The above comments for achieving a credible prediction capability for contained bursts were concerned with the situation where the key pretest limitation is that the particular test under consideration is not a part of the technology base on which the prediction procedure is based. In addition to this constraint the following are further restrictions placed on solving the ultimate detection problem:

(a) Limited knowledge of the geology - Estimates of the general mineralogy and chemical properties, density, porosity, and seismic velocity will probably exist.

(b) Almost no information on detailed material properties - Hugoniot data for the specific geology is practically never known and other laboratory data is seldom available.

(c) No information of the precise source configuration.

Bearing these obstacles in mind, along with the previously identified problems associated with defining in situ material properties, I shall attempt to suggest how the finite-difference methods (codes) can be used to aid in the solution of the key problem. Such a program must provide a sufficient data (theoretical and experimental) base to allow a synthesis of the limited data available from remote sites with the detailed data from known sites. It must isolate the key parameters that influence the late-time earth motion from underground explosions, and then provide a means of estimating these key parameters from the limited data mentioned.

For example, there are reasons to expect that one may be able to estimate Hugoniot information based on limited knowledge of the constituents of a given material (5-7). On the other hand, we do not know the degree of accuracy to which the Hugoniot information must actually be known to provide a sufficient input for theoretical calculations of the decoupling problem. In this case, a series of theoretical calculations would be most useful in establishing the requirements for obtaining Hugoniot information beyond the current state of the art.*

Similarly, parametric studies should be used to establish the key parameters and their range of effects to drive experimental work for lower stress levels down to and including the so-called "elastic

---

*Reference 3 includes some parametric studies involving variations in the source characteristics assumed in Piledriver calculations which suggest that uncertainties in the high pressure equation of state may not be as important as the material properties at lower stress levels.

regime." Such studies can evaluate general classes of material behavior (brittle, elastic-plastic, rate dependent, porous, etc) and provide the base line to identify meaningful field and laboratory experiments to establish a credible constitutive relation for in situ geologic material behavior. Personally, I believe our greatest folly in past applications of the theoretical methods has been the neglect of such systematic parametric studies in favor of attempting to predict specific events. In the remaining sections I shall return to this point with specific examples from rather recent work.

Philosophical Comments

To point out the possible sources of error in computing ground shock with finite-difference procedures, it is perhaps useful to first review the basic logic applied to solve transient continuum mechanics problems. Classical continuum mechanics theory requires that three conservation principles be obeyed, i.e., conservation of mass, momentum, and energy. These principles are usually mathematically stated in the form of integral equations, or in a more restricted sense, partial differential equations. An additional relation, a constitutive relation, is required to relate the kinematic behavior of the continuum to its thermodynamic properties at any interior point at any time. Furthermore, consistent initial and boundary conditions must be specified to assure that the mathematical problem is properly posed.

The resulting mathematical problem is usually highly complex, being nonlinear in both its kinematical and constitutive relations. Finite difference techniques have the advantage of allowing one to study this nonlinear behavior without the injection of linearizing assumptions. However, instead of obtaining exact solutions for approximate mathematical models, they yield approximate solutions of the more exact nonlinear models of reality. Thus, when one compares numbers generated by the finite-difference techniques to experimental data, he must consider two distinct sources of error.* As in the linear case, a possible source of error is the mathematical model of the physics, i.e., the constitutive relation and the initial and boundary conditions. Unlike the linear case, a second source of error is the procedure of generating the numbers itself, i.e., purely numerical errors. Too often, this second source of error has been ignored, and direct comparison between computer results and experimental data has been made without verification that such a comparison is valid.

Valid comparisons between finite-difference results and experimental results can only be made after both numerical and experimental errors have been evaluated. In emphasizing this point, an analogy may be instructive. If an experimenter wishes to study some

---

*Or three, if one considers the experimental error as well.

phenomena in the laboratory, he determines the required instrumentation
sensitivity, obtains such instrumentation, calibrates that instrumenta-
tion against some standard, and then makes his measurements. Such a
procedure allows him to state the accuracy of his measurements. The
finite-difference techniques represent both the phenomena to be studied
and the instrumentation which makes the measurements. The physics is
represented by finite-difference analogs of the conservation laws and
the assumed constitutive relation, and the whole calculational process
provides the numbers analogous to a measuring device. Hence, the usual
experimental calibration process is analogous to the determination of
the accuracy of a given numerical procedure, including the representa-
tion of the physics.

Numerical errors have two primary sources. The first is
introduced when a continuum is represented by a discrete number of
points. The second source is the process of solution of the discrete
system. In considering the first source, it is instructive to remember
that the equations of continuum mechanics can be derived by applying
the conservation principles to small discrete sections of the continuum,
writing difference equations, ignoring higher than first order terms,
and taking the limit as the discrete section shrinks to a point in the
space-time domain. Thus, the primary difference between the difference
equations and the differential or integral equations of continuum
mechanics is the absence of the limiting procedure and the possible
omission of important second order terms. It is also clear that as
smaller and smaller sections of the continuum are taken (finer and
finer zoning), the difference equations approach the proper differ-
ential equations. (Note that this does not necessarily mean that the
solution of the difference equations converges to the solution of the
limiting differential equations. Although we proceed with confidence
in the finite-difference techniques, no theorem stating this desired
result has been proven except in very special cases.)

Associated with this discussion of the limiting process, it
is worth noting that most finite-difference equations have been derived
in a mathematical sense by writing finite-difference analogs for partial
differential equations with little regard for the physics involved. In
so doing, it appears that one has a number of arbitrary choices he can
make in prescribing the actual form of the difference equations that
all approach the proper differential equations in the infinitesimal
limit. Because the difference equations all approach the correct
limiting equations, it is reasonable to ask if they are all equally
reliable in their finite-difference form or if one is better than the
others. Few definitive sutdies addressing this problem have actually
been reported. To my own personal knowledge, actual comparative studies
suggest that difference methods based on physical insight which to some
close degree satisfy the conservation principles in the explicit finite-
difference form are preferable. If one starts with the more basic
physics and requires a strict adherence to the conservation principles
in the finite-difference formulation associated with a finite piece of
matter, then the arbitrariness in writing difference equations is

63

significantly reduced. In fact, it is a nontrivial task to find a set of internally consistent difference equations (8-11).

Usually the various codes strictly conserve momentum analogs, but the ways in which they conserve energy and mass have subtle differences. Some conserve total energy by definition. They compute a kinetic energy analog and subtract it from the total energy obtained from a total energy transport analog to obtain an internal energy quantity to be used in the constitutive relation. Other codes compute changes in both kinetic and internal energy analogs and check each time step to be certain that total energy is conserved to within one part in a large number, say 10^6. Some use kinetic and internal energy analogs defined in such a way that the finite difference equations explicitly conserve total energy. In the infinitesimal limit, all of the codes reduce to the correct form of an energy conservation principle.

The definition of mass conservation is simple since it merely requires that whatever flows out of a given zone must flow into one of the adjacent zones. Hence, mass conservation in a Eulerian calculation is simply a bookkeeping job. The heart of the matter comes with the manner in which mass transport is calculated and centers around the question of what density is to be assigned to material that flows across a zone boundary during a time step. Depending on this definition, the transport may be calculated with backward, forward, or centered difference equations. The centered differencing scheme is correct to second order whereas the other two are first order schemes. Since the effect of transport is to "smear" a pulse, the first order schemes provide more "numerical" diffusion than does the centered procedure. In fact, no artificial viscosity* (12) is required if first order difference equations are used for transport. This might appear to be an advantage for first order systems, but it must be pointed out that they introduce an uncontrollable numerical linear viscosity (13) into the calculation which is as unreal as the artificial quadratic viscosity. Further, a linear viscosity erodes a pulse much more rapidly than does a quadratic viscosity (10). It has been found that the use of centered differencing and an artificial quadratic viscosity produce peak pressures about 20 to 30 percent higher than backward differenced transport (14).

As was previously pointed out, one hopes that the correct solution for a given problem is approached as finer and finer zoning is employed. With a one-dimensional code and current computer capabilities, it is possible to use "sufficiently" fine zoning for most problems. However, it is not practical to zone a two-dimensional problem sufficiently fine throughout. Hence, most two-dimensional codes employ rezoning techniques such that a given number of zones are employed over a small region surrounding the source and then rezoning

---

*Richtmeyer and von Neuman originally introduced the artificial quadratic viscosity in order to calculate problems involving strong shocks. Otherwise large oscillations are introduced when the difference equations are applied to a region in the neighborhood of a discontinuous function.

into larger dimensions as the problem progresses. A typical way of rezoning in a two-dimensional code is to replace four small zones by a single large zone that has dimensions twice as large as each of the small zones.

A problem which must be solved in the "dimension-doubling" technique of rezoning is the manner in which the pressure and velocity are assigned to the new larger zones. The mass of the new zone is simply the sum of the masses in the four smaller zones. The new momentum can be similarly defined. When this is done a velocity for the new zone can be obtained by dividing the new momentum by the new mass. This velocity can then be used to compute the kinetic energies in the four old zones. Hence, if total energy is conserved by definition, we see that a different partition of energy between kinetic and internal energy is caused by the rezone. The calculation then proceeds, and one hopes that no great error has been introduced.

Summarizing, in considering the errors associated with finite difference techniques, there are three questions to be answered:

(1) Are the numbers generated an accurate solution of the finite-difference equations?

(2) Is the real solution of the difference equations an accurate solution of the posed continuum mechanics problem?

(3) Is the posed continuum problem, including the constitutive relation and initial and boundary conditions, an accurate model of reality?

Ultimately, these questions must be independently answered or we are destined to obtain right answers for the wrong reasons or wrong answers for the right reasons--both unacceptable results from a scientific point of view.

Some Examples of Numerical Errors

This section will review several sources of numerical errors and will provide a couple of examples suggesting that one should always be a bit skeptical of code-generated numbers until he satisfies himself of their validity by considerable study of the specific problem being calculated. Before dealing with specific cases, some general comments are probably in order.

Earlier, the question was raised as to whether or not one scheme of differencing is better than another. The answer to the question probably is that it depends on the specific problem being considered. There are, however, some points to make concerning the inadequacies common to most, if not all, finite-difference methods. For example, one may deceive himself as to the accuracy he actually achieves

65

by applying high order differencing schemes. Several years ago, Trulio (10) showed that the maximum rate of convergence of the Richtmeyer-von Neuman difference equations for pulse propagation in a linearly elastic material (which are second order based on a Taylor expansion error analysis) is proportional to $(\Delta X)^{3/2}$. Moreover, the result of his study showed that convergence is not uniform and that decreasing the zone size could in fact lead to less accurate answers in some cases.

More recently, it is my understanding that workers at the California Institute of Technology have more rigorously demonstrated that higher than first-order Eulerian difference methods for multi-dimensional hyperbolic problems do not give higher accuracy (15). Intuitively this is to be expected because the error term in the Taylor series includes a derivative of some field variable which does not exist at shock fronts.

Another source of error in the application of finite-difference calculations results from the application of an artificial viscosity to reduce numerical oscillations near shock fronts. Generally, the use of this artifice reduces peak values--particularly the application of a linear rather than a quadratic viscosity relation. This source of error, as well as others, will be commented on further in the following examples.

PANCAKE Problems

A 1966 study (16) attempted to provide justification for emphasizing research in addressing questions of numerical accuracy. Actually, it focused on consistency rather than accuracy, and posed the question of whether or not different finite-difference procedures would produce the same answer for the same defined physical problem. Obviously, if consistency is lacking, then the question of accuracy cannot be avoided.

Calculations were made on the AFTON code (11), the OIL code (17), the PIC code (18), and the FLU code (19). These four codes were independently developed under DASA and AFWL contracts and were used in the studies reported in Reference 16 by personnel from the organizations responsible for their development. The problems computed were specified such that all of the participants used the same constitutive relation, the same source definition, and the same initial zoning. Hence, any differences in calculated numbers could not be related to physical un-certainties, i.e., their source was entirely numerical.

The initial geometry of the problems involved a thin disc source of energy placed at an air-ground interface such that the top of the disc was flush with the interface as shown in Figure 21. (Because of the initial geometry, this series of calculations has become known as the PANCAKE Calculations.) The pancake was initially 2-cm thick and 14.8 cm in radius.
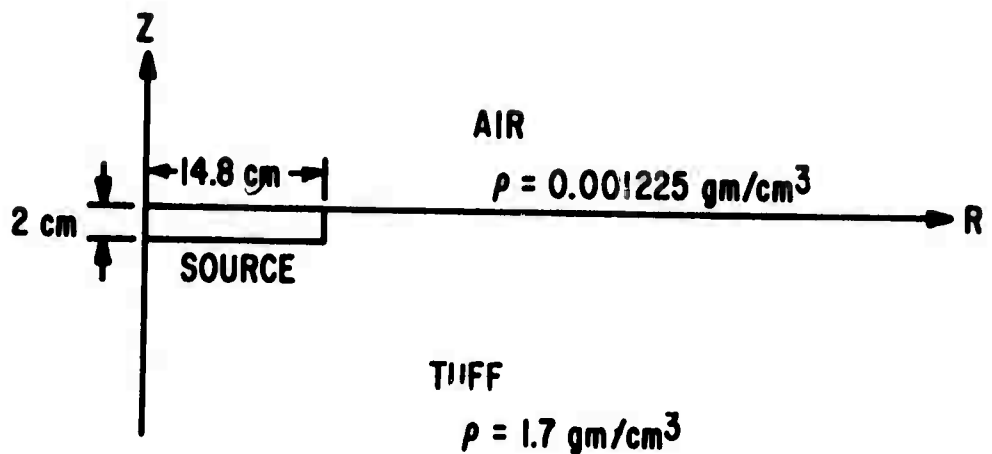
Figure 21. PANCAKE Problem Geometry.

$$P = \left[a + \frac{b}{g}\right] \rho E + A\mu + B\mu^2 \qquad \eta \geq 1 \text{ or } (E < E_s \text{ and } \eta > 1/V_s)$$

$$= a\rho E + \frac{b\rho E}{g} + A\mu \exp(-\beta h) \exp(-\alpha h^2) \qquad \text{otherwise}$$

where $\eta = \frac{\rho}{\rho_0}$

$h = -\frac{\mu}{\eta}$

$\mu = 1$

$g = 1 + \frac{E}{E_0 \eta^2}$

Here, P is the pressure in megabars, ρ is the density in g/cm³, E is the internal energy density in megabar-cm³/g, and ρ_0 is the initial density. The various parameters for the ground (tuff) were defined as

a = 0.5	A = 0.064 Mb	E_0 = 0.005 Mb-cm³/g
b = 1.1	B = 0.07 Mb	E_s = 0.10 Mb-cm³/g
V_s = 1000	ρ_0 = 1.7 g/cm³	$\alpha = \beta$ = 5.0

These same parameters were used for the constitutive equation for the air except that ρ_0 = 0.001225 g/cm³ in the air.

The number of zones in the problem were 66 vertical and 46 radial zones (3036 zones total) such that 40 vertical zones were in the tuff. Initially, all zones were square, 0.4 cm on a side so that there were 185 zones in the pancake source (five zones thick and 37 zones in radius).

The computations were started by a uniform deposition of 200 tons (HE equivalent) of internal energy (837.2 jerks) into the pancake at zero hydrodynamic time.

Figures 22-24 show the calculated peak pressure on-axis (R = 0) beneath the source, along a 45-deg line and at the air-tuff interface. The spread in peak pressure is on the order of a factor of two. The PIC results are consistently higher beneath the source, along the 45-deg radial, and along the air-tuff interface. Because it is felt that the results of all of the codes were generally low, the PIC results are probably the most nearly correct. This is not intended as a blanket endorsement of PIC because the number of particles involved complicates anything but a direct comparison of the numbers obtained in the calculation. (The code accuracy increases with both increasing density of zones and particles. In this particular calculation, 20 particles per initial zone were used.) Ultimately, the best code is the one which obtains the most accurate numbers for a fixed cost, or conversely, a given accuracy for the least cost.

In Figure 22, note the departure of the OIL and FLU results at approximately 23-cm depth. This point approximately corresponds to the point at which a rezone occurred in the OIL calculation. The kink in the PIC results at a range of 6 cm also appears to be associated with a rezone. The AFTON results are smooth, probably because the problem was continuously rather than discretely rezoned. (When grid motion is initiated or stopped abruptly in AFTON, kinks generally appear (21).)

In Figure 23, note the radical separation of the AFTON results beginning at a range of about 40 cm. This result is probably due to the fact that the "accordion" coordinate system began to significantly outrun the shock on the 45-deg radial, thereby smearing the momentum over larger effective zone distances and producing lower pressures. Again the spread in inconsistency is in excess of a factor of two for ranges greater than about 50 cm. A similar result is found for the pressures a half-zone below the air-tuff interface shown in Figure 24.

Note the kinks in the AFTON, OIL, and PIC on-axis particle velocity results in Figure 25. The origin of these kinks is unknown, and is made more of a mystery since they do not reflect into the pressure attenuation curves shown in Figure 22. It is suspected that they are unreal. The hook in the AFTON results at 60-cm range is definitely not real and can be related to the fact that reflections from the edge of the finite-difference mesh occurred.

68
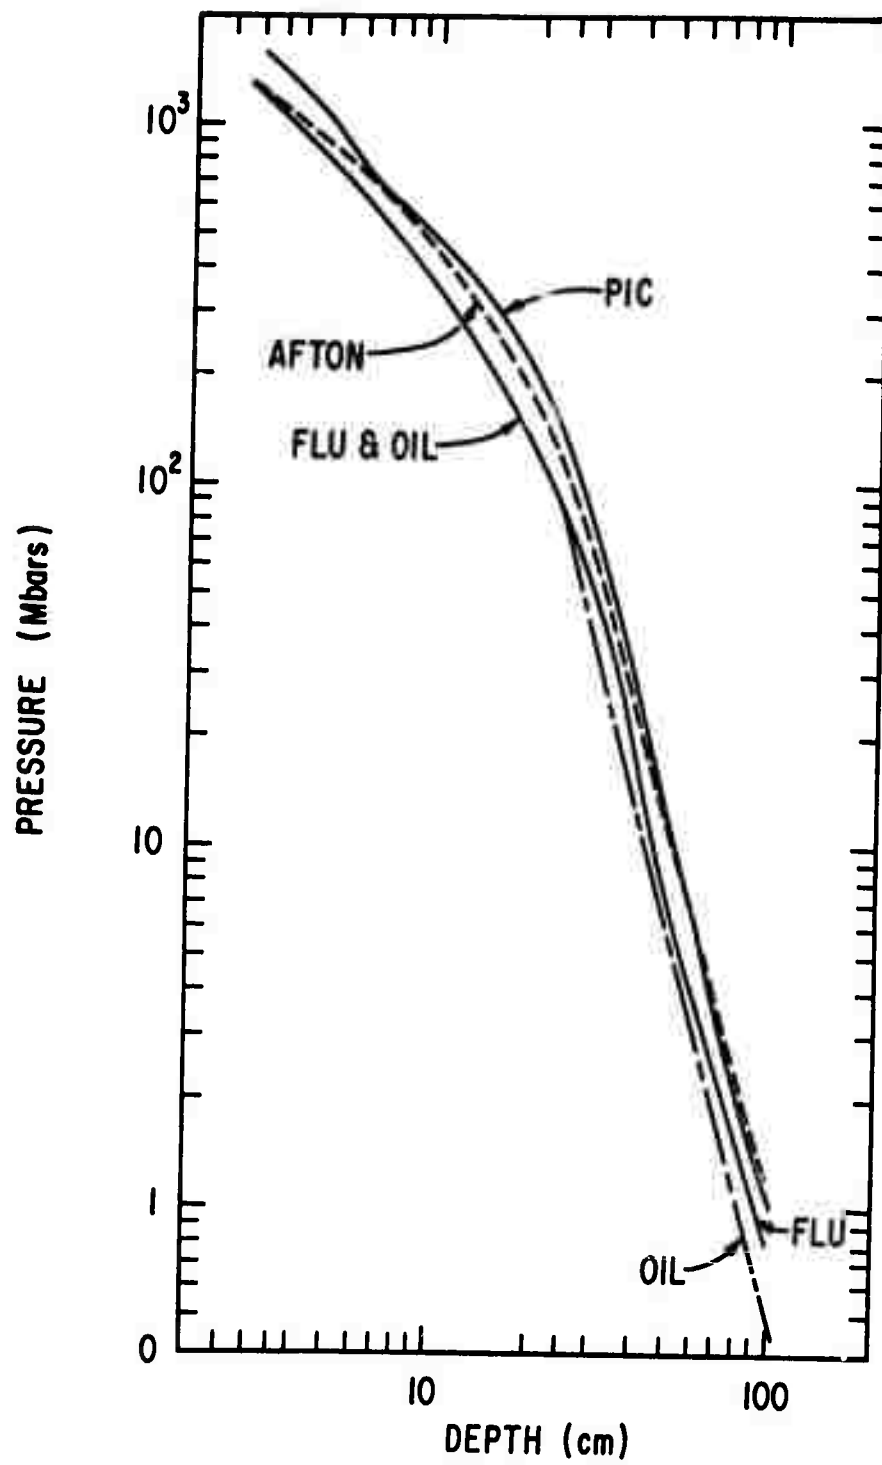
Figure 22. Peak Pressure On-Axis. R = 0.
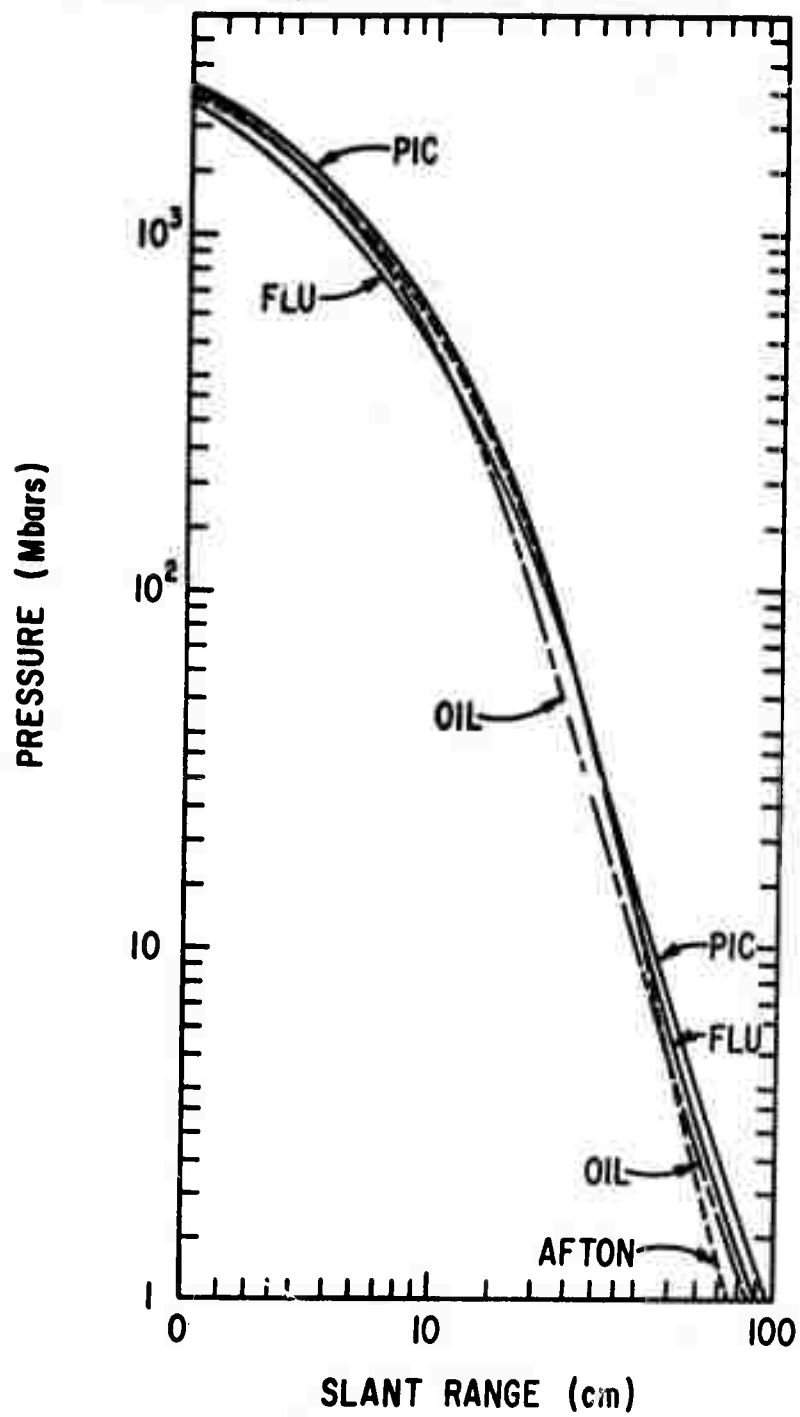
Figure 23. Peak Pressure on the 45-deg Radial.
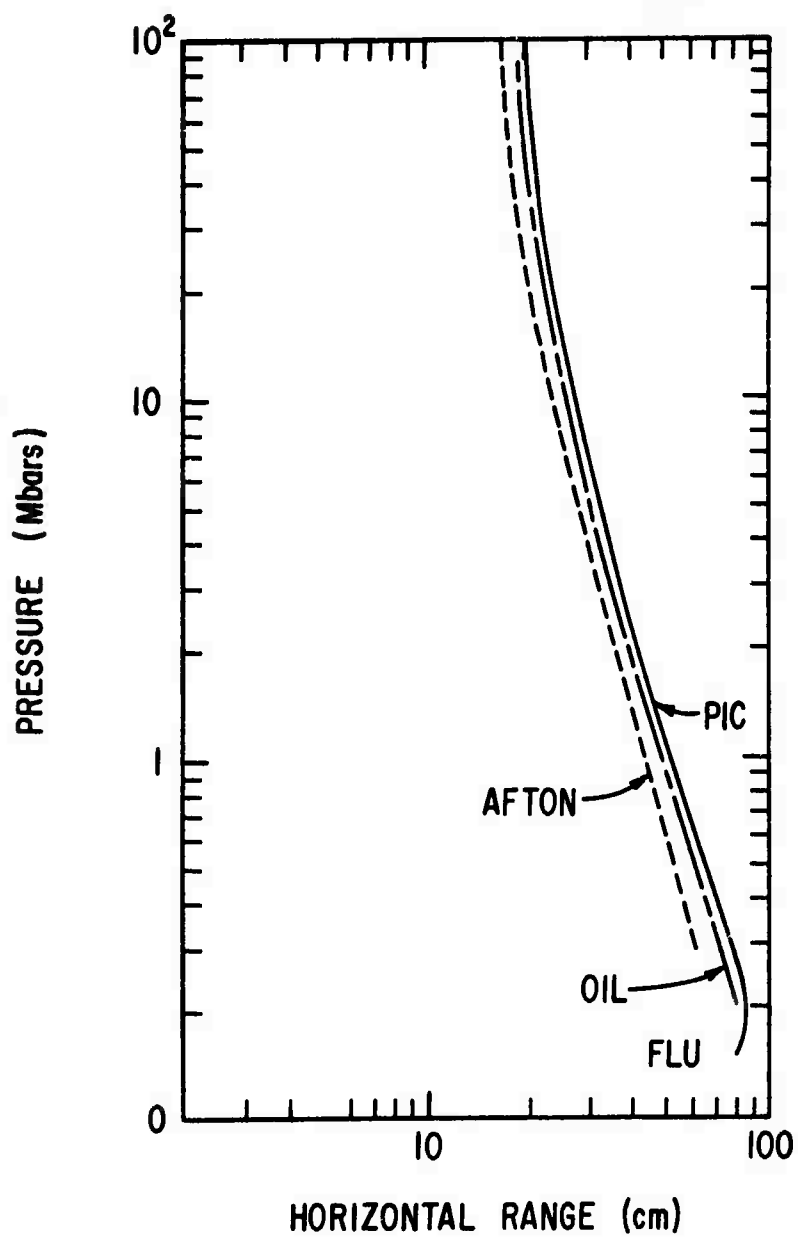
Figure 24. Peak Pressure at the Air-Tuff Interface.
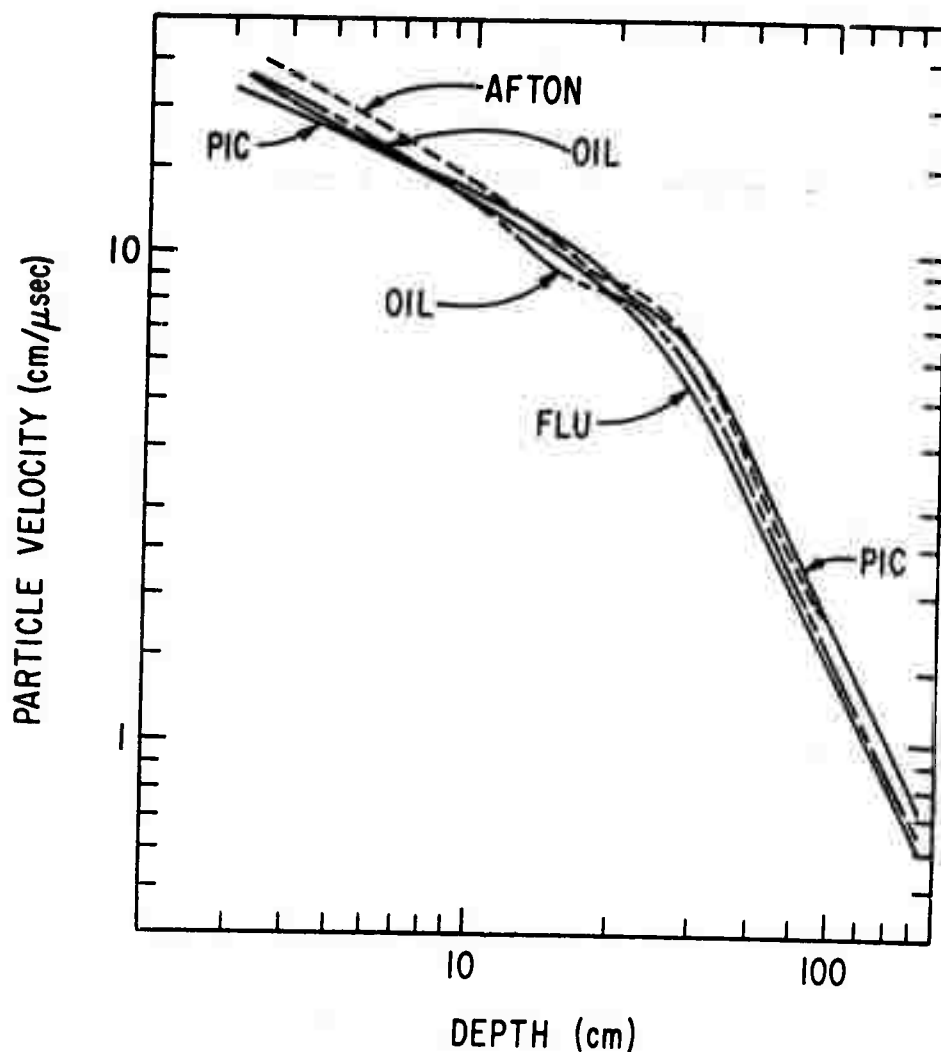
Figure 25. Peak Vertical Particle Velocity On-Axis. R = 0.

In general, it was noted in Reference 16 that significant inconsistencies were often larger than a factor of two (ratio of highest to lowest). The agreement between peak particle velocity attenuation plots was better. For the case in question (strong shocks), one would expect the pressure to be proportional to the square of the velocity. However, the particle velocity results of this study agreed better than one would predict based on this expectation.

The peak value differences between the AFTON, OIL, and FLU results were generally within three zone-widths, which may define some consistency uncertainty. However, the relevant question to be answered is whether or not that is an acceptable limit. In the case of large zones and rapid attenuation rates, three zone-widths may correspond to large peak value differences.

72

The various spatial profiles of pressure and particle velocity at various times demonstrated more prominent differences than did the peak value plots. For example, consider the difference in pulse widths, area under the pulses, and the detailed kinks in the pulses in Figures 26-28. Although there is not a direct conversion, these same differences would also be reflected in the time histories at various specified points. Which do you believe?

What causes the differences? In a rather detailed study of zoning and rezoning procedures, Trulio (21) showed that kinks and detailed changes in peak value attenuation curves are originated in AFTON when one changes the motion of the finite-difference coordinate mesh in a continuous way. It seems reasonable, and indeed may be proven, that abrupt changes (such as discrete rezoning) will also introduce kinks. The pertinent and, at present, relatively unanswered question is, how big are the kinks? It is clear, at least to the author, that they are large enough for one to be skeptical of accrediting equation of state or physical details with the origin of many of the various perturbations that generally occur in the results from calculations of this type.

The assertion that the inconsistencies result primarily from zoning and rezoning details was further substantiated by other calculations of the problems posed in Reference 16. An OIL calculation with 0.2 x 0.2-cm zoning (instead of the 0.4 x 0.4-cm zoning used in the study) gave results generally in agreement with the AFTON results on axis to a distance of 40 cm. Figure 29 gives the on-axis results for the internal energy problem for initial zoning of 2 x 2 cm, 0.4 x 0.4 cm, and 0.2 x 0.6 cm. Note that the rectangular 0.2 x 0.6-cm zoning gives results that are about 1.3 higher than those for this study beyond a depth of 60 cm and that the AFTON results were the highest of the four codes applied in the study for this range.

To further complicate matters, Lagrangian calculations gave results that were different from those calculated by Eulerian procedures. In fact, a Lagrangian calculation--with fewer zones than those in this study--on MOTET gave results consistent with and, in some cases, slightly higher than the AFTON results on-axis out to about 40 cm. Also, a later coupled calculation with ELK showed that the introduction of a Lagrangian mesh increased the on-axis pressures. The implication is that mass transport calculations tend to degrade pulses, i.e., more zones are required to obtain a given accuracy with an Eulerian calculation than with a Lagrangian calculation. It should be noted, however, that Lagrangian calculations tend to overshoot peak values. In any case, it is usually expected that the calculated numbers become more accurate as more zones are introduced in the problem. (However, note the earlier comments suggesting that the convergence is not uniform.)

Thus far, the discussion has dealt primarily with consistency. It has been implied that all of the calculated peak values considered in Reference 16 were generally low, and that the correct answer would be

73
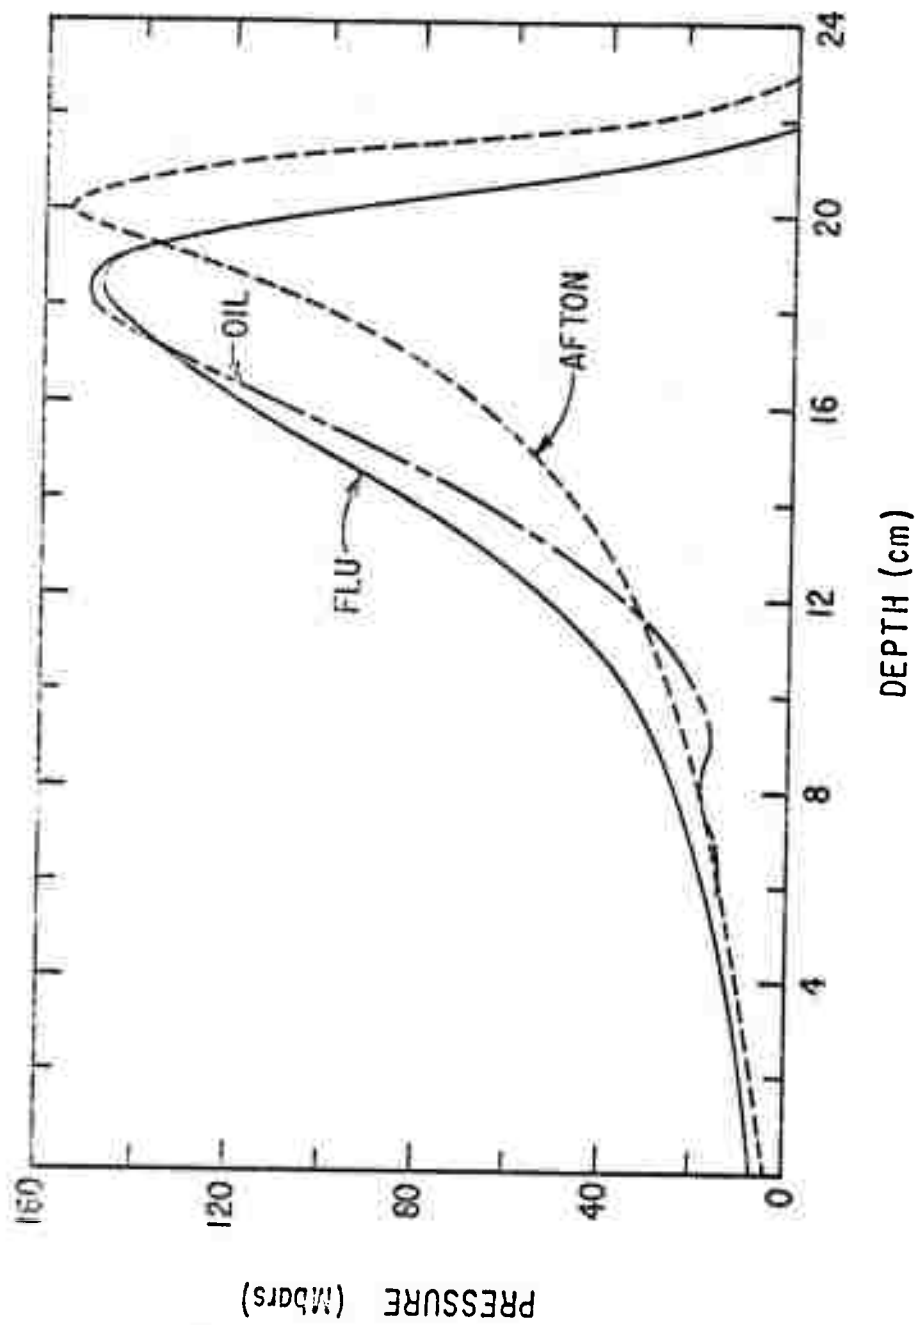
Figure 26. Pressure Profile On-Axis at t = 1 μsec.

74
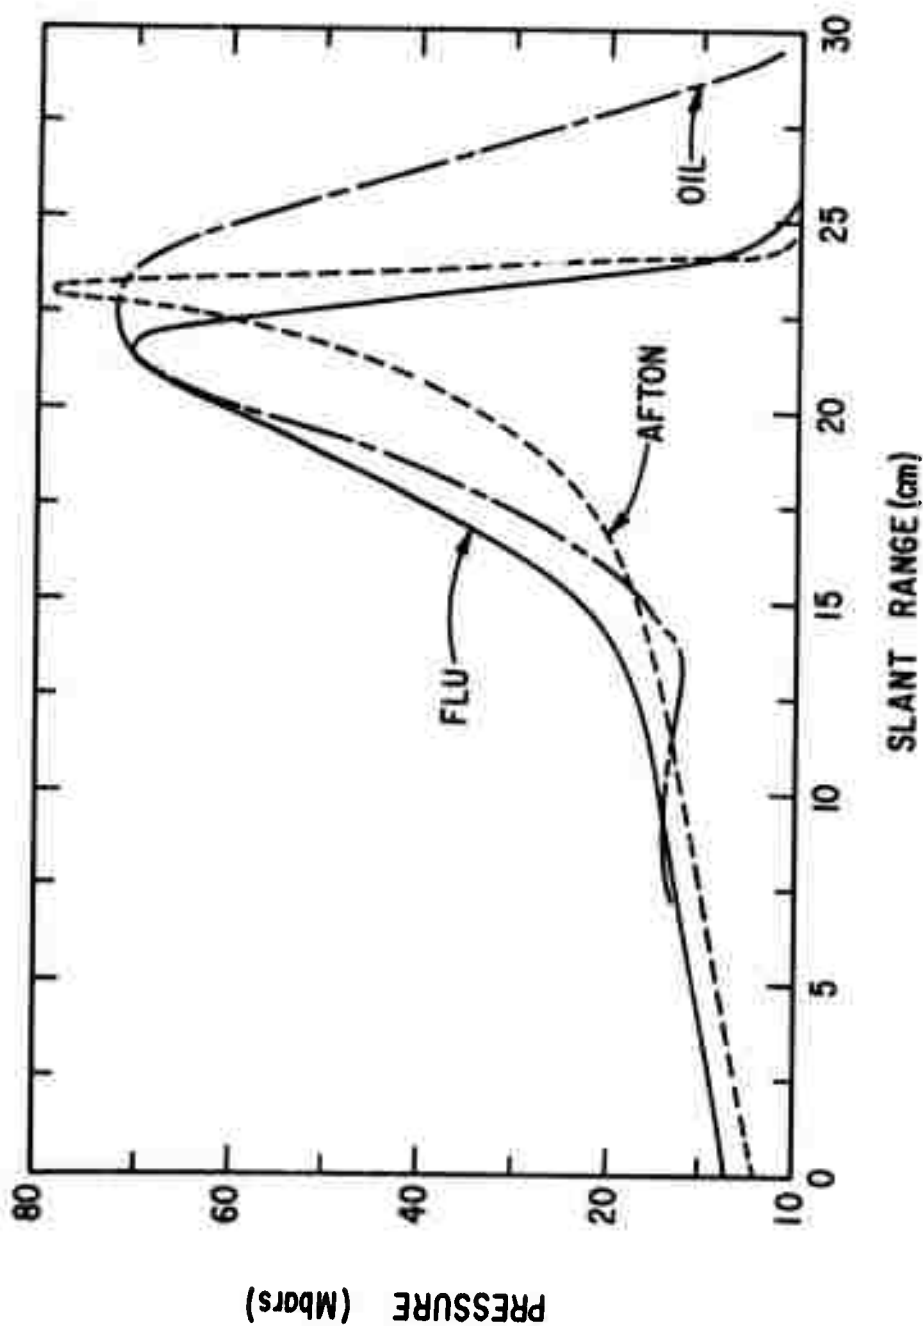
Figure 27.   Pressure Profile on the 45° Radial at t = 1 μsec.
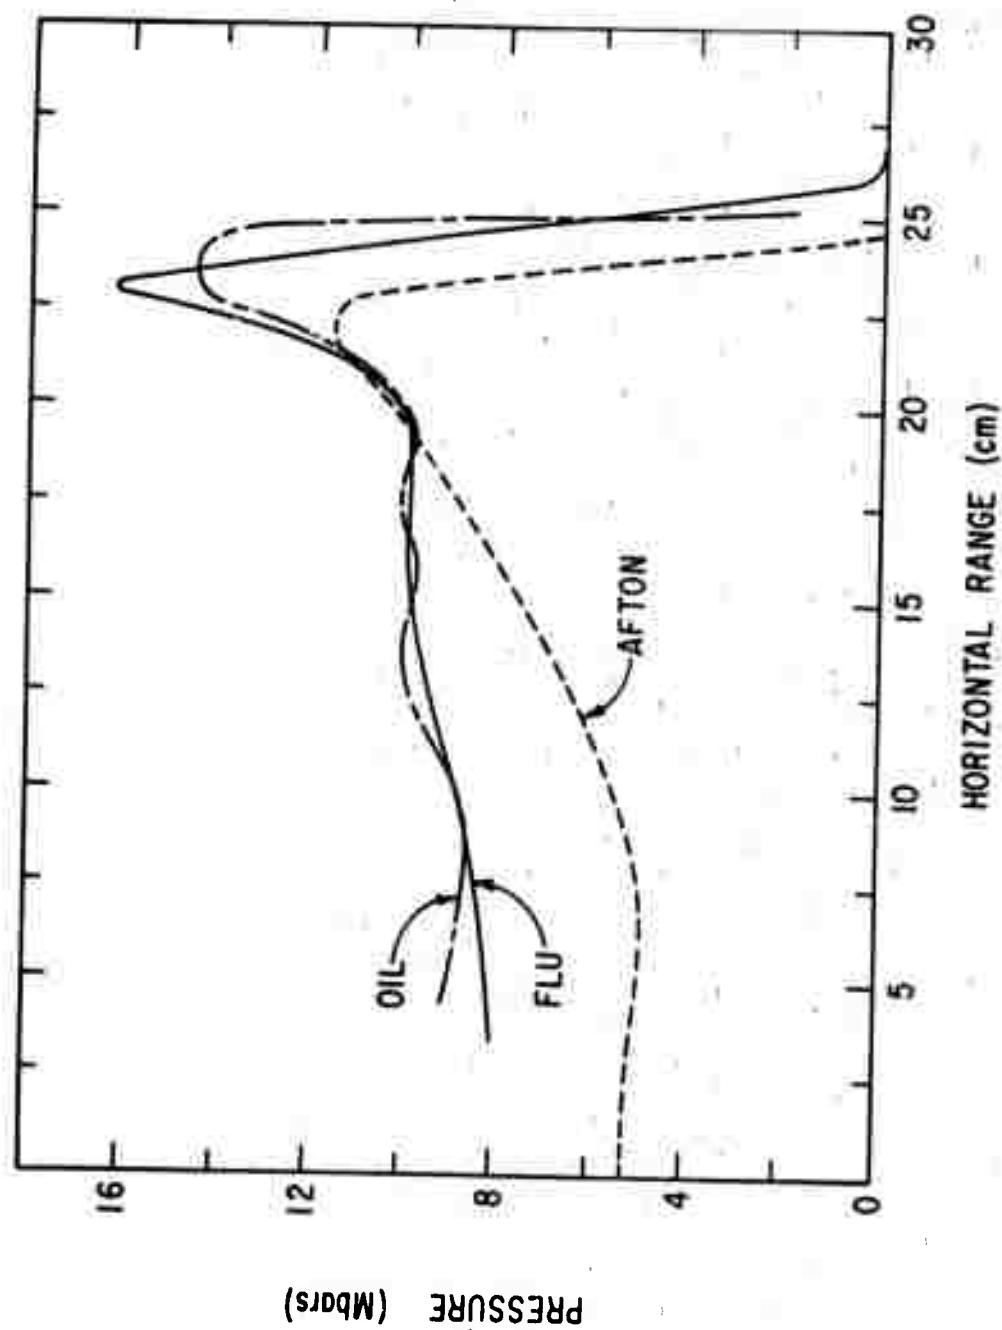
Figure 28. Pressure Profile at the Air-Tuff Interface.

76
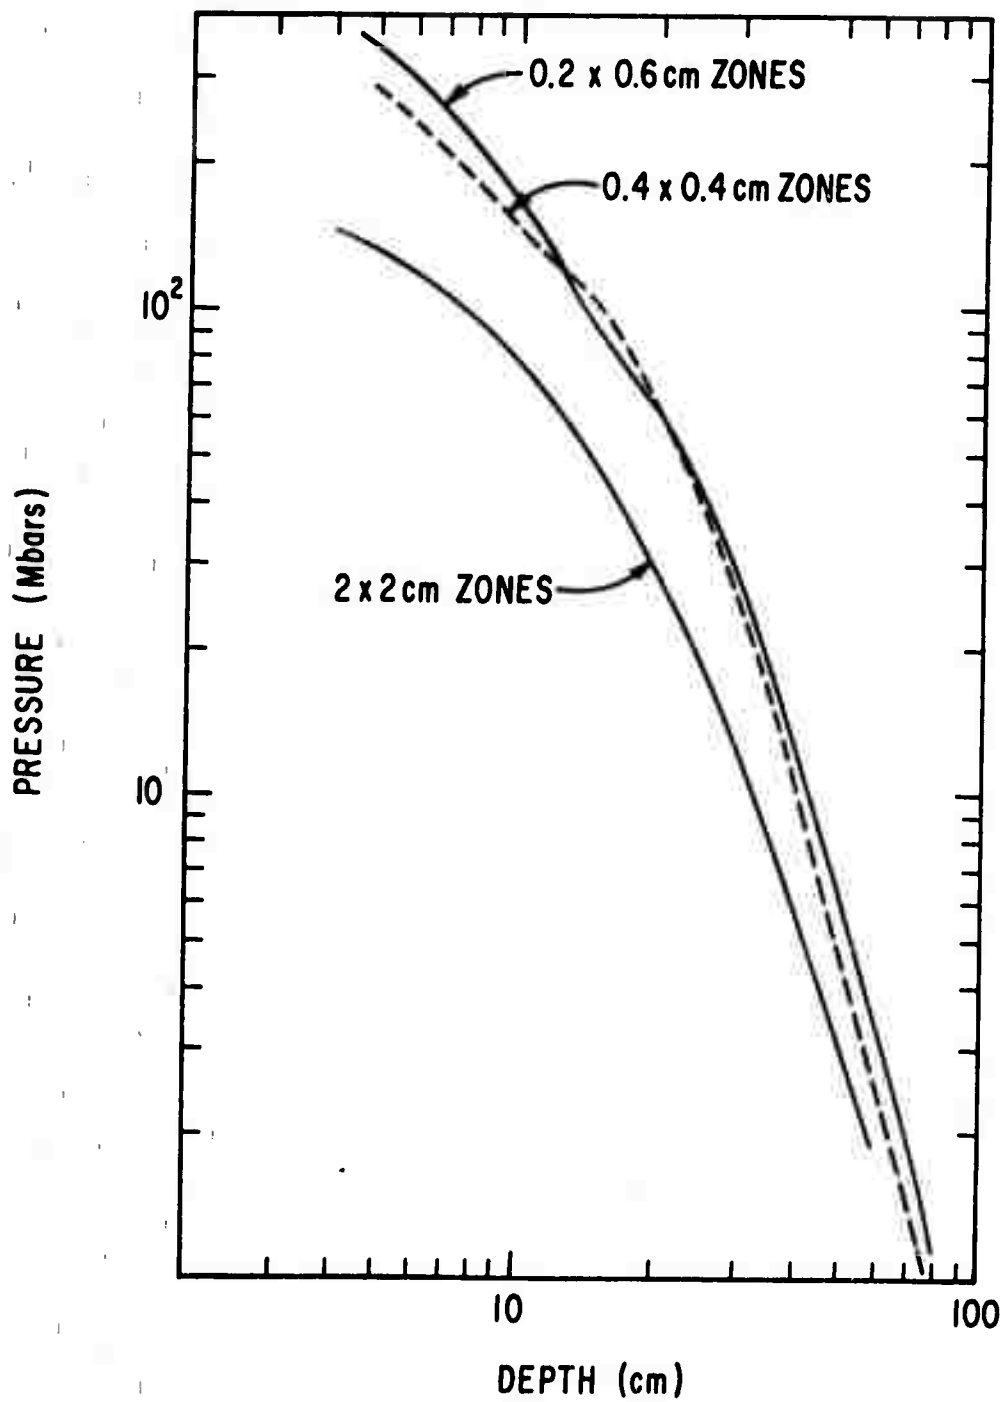
Figure 29. Effect of Changing Initial Zoning on the Peak Pressure
On-Axis. R = 0.

approached with finer and finer zoning. In order to estimate the accuracy of the calculations on-axis, a very finely zoned, one-dimensional, plane Lagrangian calculation with no rezoning was run. The solution for these problems should be identical to those for the two-dimensional problems on-axis until a rarefaction from the edge of the pancake reaches the axis (about 20-cm depth). Beyond this point, no direct comparison between the one- and two-dimensional problems can be made. To obtain some feeling for the error which might exist, a one-dimensional calculation was run on AFTON with the same zoning as the on-axis zoning of the two-dimensional problem. In addition, the problem was "rezoned" in the same way as in the two-dimensional calculation. The results of this calculation are shown in Figure 30. The error in the coarser-zoned, one-dimensional calculation is about 30 percent (assuming that the finely-zoned results are correct). It is difficult to imagine how the two-dimensional calculations could be in less error than this one-dimensional result. It can be argued that one-dimensional estimates are valid for the errors that propagate on-axis. Errors associated with the transport of momentum away from the axis could degrade peaks even further.

Recent studies of the convergence characteristics of the PANCAKE problems suggest that the errors in FLU were fairly independent of direction, however (22).

In any case, we arrive at the conclusion that the study of numerical errors is quite a pertinent subject; they should not be ignored. Any calculation performed should include subcalculations that would serve to estimate some bounds on the error involved. In particular, zoning and rezoning features should be scrutinized.

Spherical Waves in an Elastic Medium

It is generally accepted that one-dimensional calculations can usually be zoned sufficiently fine to obtain accurate numerical solutions. In general, this may be the case, but there are ways that one can obtain inaccurate results even with fine zones. One such example is given in Reference 23 where comparisons are made between code results and an exact analytical solution for spherical wave propagation in an elastic medium. The problem geometry is shown in Figure 31.

At the time this problem was solved, its primary motivation was to serve as a check case for several one-dimensional, finite-difference techniques developed under DASA sponsorship and AFWL technical direction. A typical starting condition used in these codes was to dump internal energy into a gamma-law gas source region. The analytical solution was obtained for the quasi-static expansion of the gamma-law gas obtained in Figure 31. The quasi-static behavior of the
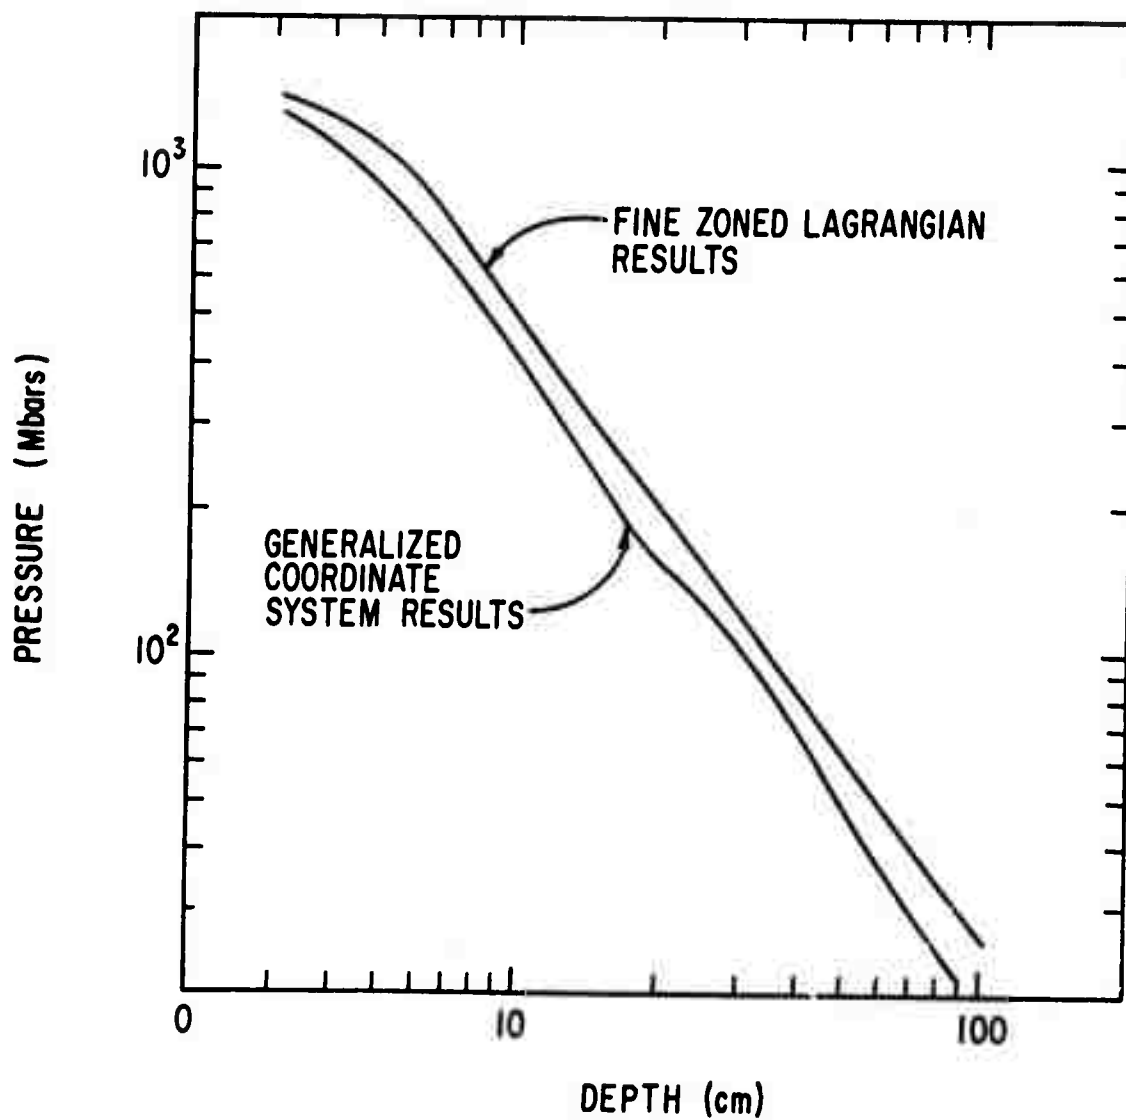
Figure 30. One-Dimensional Calculation of the PANCAKE Problem.

79
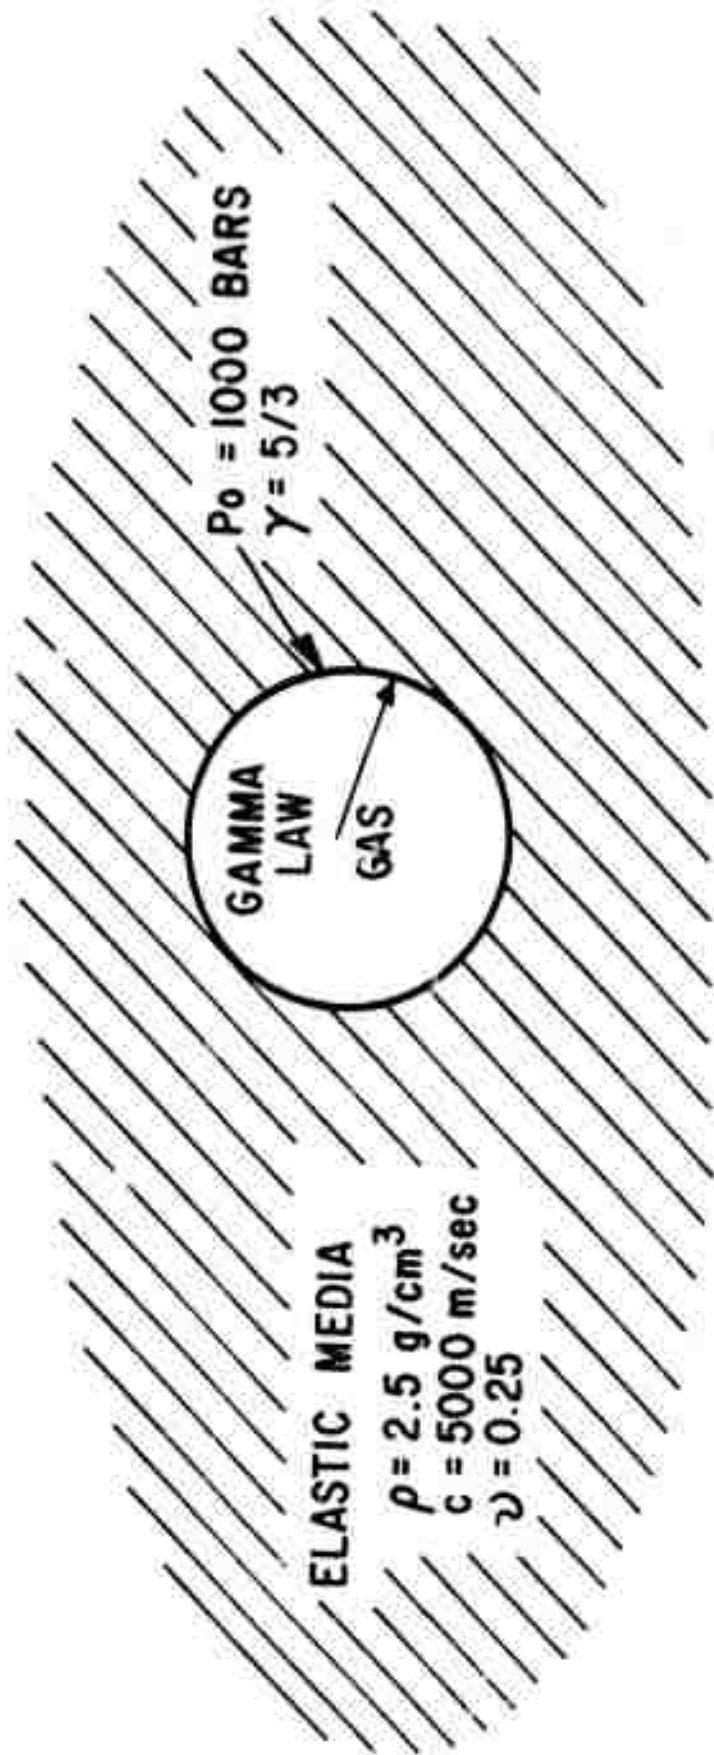
Figure 31. Problem Geometry for Spherical Wave Studies.

ELASTIC MEDIA

$\rho = 2.5$ g/cm^3
$c = 5000$ m/sec
$\upsilon = 0.25$

Po = 1000 BARS
$\gamma = 5/3$

GAMMA
LAW

GAS

gas in the finite-difference calculation was obtained by using a single zone to represent the cavity region, thus assuring "average" behavior*.

Before the results are presented it may be of interest to review what motivated the numerical study. Initially, the problem was of interest purely as a straightforward check case. The first numerical results that were obtained were badly in error, both in pulse shape and attenuation phenomena. Further, no amount of refinement of the finite-difference mesh, i.e., fine zoning, in the elastic region improved the result. In fact, beyond some point, finer zoning in the elastic medium produced even greater errors. Efforts were made to improve various finite-difference approximations within the code with no improvement of the results.

In all of the calculations made up to this time, no great care had been taken in prescribing the density of the gas in the cavity, because the theoretical result indicated that the solution would be independent of density variations for the geometry and pressure level treated. At this point, several numerical calculations that varied the initial properties of the gas were made. The results indicated that the numerical solution was indeed dependent on properties of the gas other than the initial pressure.

These results were due to a numerical phenomenon, associated with many finite-difference techniques, which should possibly have been expected at the outset. Most of the finite-difference techniques do not completely satisfy the physical boundary conditions at an interface, i.e., displacements are continuous across the interface but the requirement that the traction be continuous across the interface is not explicitly used in the calculation. In fact, in most codes, the boundary zones are treated no differently than any other zone. The momentum at a mesh point is usually calculated by multiplying the velocity of a grid line by the sum of the mass contained in the "half-zones" adjacent to the grid line. Hence, in the problem treated here, the momentum associated with the grid line that defined the cavity wall is dependent on the density of the gas contained in the cavity. There is possibly some density that may be defined, which will result in the proper boundary condition, namely, that the pressure in the cavity be equal to the radial stress in the elastic medium at the cavity wall.

Figure 32 gives the computer results for particle velocity profile at a time of 10^{-5} sec, at which time the pulse has propagated five cavity radii into the elastic medium. The mass ratio indicated on these graphs is the ratio of the mass in the half zone of gas adjacent to the cavity wall to that of the half zone of "rock" adjacent to the cavity wall.

---

*The calculations were performed on POD, a typical one-dimensional Lagrangian code developed by Physics International (24).
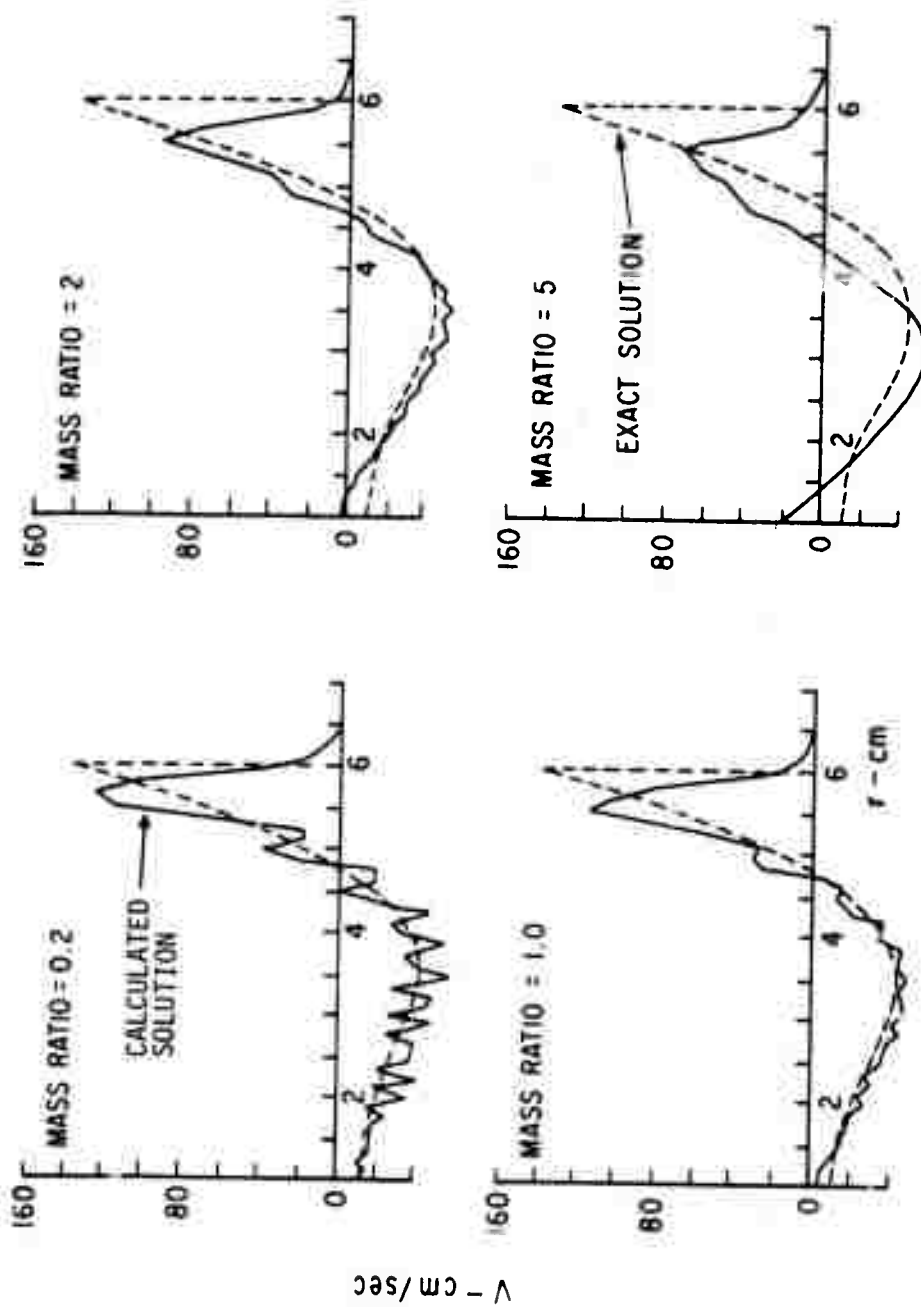
Figure 32. Effect of Boundary Condition on Particle Velocity Profile. Solid line is calculated solution; dotted line is exact solution.

82

Note that the results become smoother and the peak lower as this ratio is increased. Note also that the profiles are "quantitatively" correct in all cases (to the point that one might assume any one of them as correct if the exact solution were not known); however, they are quantitatively in error. Note also that as the results become smoother, the error increases.

Figure 33 presents peak particle velocity versus radius plots. Note that as the mass ratio is increased, the error in peak particle velocity increases. However, the error in attenuation rate decreases with increasing mass ratio.

In the above problems, there were eight zones per cm; thus, there were eight zones in the theoretically predicted pulse width (there are about 12 zones in the numerical pulse width). Other runs, involving 4, 16, and 32 zones per cm, were made. With the exception of four zones per cm, those results were generally the same as in Figures 32 and 33. (Peaks were slightly higher, and oscillations were smaller.)

The problem treated here was solved exactly and compared with the numerical results. This comparison suggests that these errors can be significant and, at the same, the results may appear to be correct (in the sense that, if no exact solution existed, they are of the expected form). Perhaps, one would argue, by intelligent choice of zone sizes, one could properly match the mass across the cavity wall. No doubt this is true for some problems, and various rules of thumb might be developed to specify the proper condition for certain classes of problems. However, in other problems (where there is a large density discontinuity), such a procedure cannot be defined because of other inconsistencies involving the transient time across a zone and the requirement for some desired zone definition. It would therefore appear that the only real solution, and in fact the only logical solution, is to require that the proper boundary conditions be satisfied in the finite-difference scheme (9). Then, it will not be necessary to be concerned about the density mismatch at an interface.

If the "proper" mismatch condition was specified in the problems treated (the one usually accepted as correct is a mass ratio of one), then we may further comment on the numerical errors. Depending on how far the pulse has traveled, the error in peak amplitude can be significant. At a distance of 10 cavity radii (which is not very far for most problems of interest) it is seen from Figure 33 that the error is about 25 percent based on the exact solution, or 30 percent based on the numerical solution (which in most cases will be the only solution). This error is a function of the number of zones placed in the pulse. The problem in many of the finite difference techniques currently used is that it is not feasible to use enough zones to give much better definition than that presented here, particularly in two dimensions. Perhaps the use of generalized coordinates as proposed by Reference 25 will solve a major portion of this problem.

83
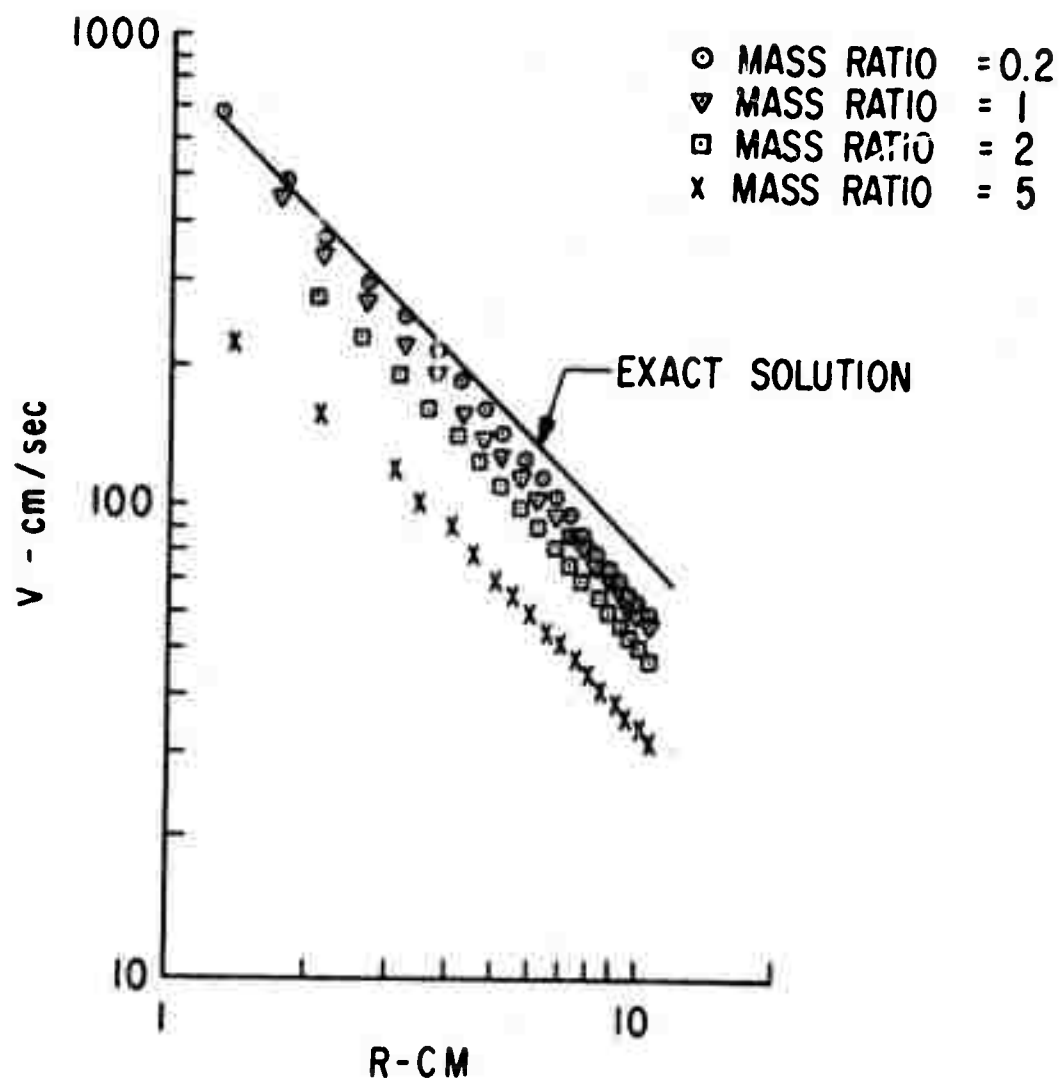
Figure 33.   Effect of Boundary Condition on Peak Particle Velocity
Attenuation.

It should be noted that this problem is most acute for calculations involving more than one dimension, especially for great distances from the source.

It may be argued that problems of the type presented here are simple while those for which the various finite-difference techniques were developed are complex. Therefore, it is most difficult to correlate information found on the simple problems with the results for the complex problems. The author would agree with this argument if the errors found for the simple problem were insignificant. However, it has been found that the errors were not insignificant, and there is no logical reason for believing that the numerical errors in more complex problems will be less.

Summary

Implicit in the above comments is the observation that direct comparisons of peak stress and velocity between calculations and experiment may not be especially revealing. This is especially true if the constitutive relation for the media under consideration is not known. In such cases, the code user is tempted to play games with free parameters until he can reproduce some particular facet of the experimental results. However, this does not, in the author's opinion, lead to increased confidence in the calculational techniques.

The "normalizing" procedure brings us no closer to the day when we can rationally set error limits for a prediction of a situation where no experiment is possible, and this must be the desired end. The discussion of numerical results in terms of their "reasonableness" must give way to the evaluation of such results in terms of their accuracy.

In order to achieve this end, the study of constitutive relations and numerical accuracy must proceed along complementary but distinct paths. In the optimum situation, constitutive relations would be independently studied by experiment. However, the codes can also be of use in this area as will be suggested in the following section.

Code Applications

The previous section presented examples of significant calculational error. While such cases are far from unique and it is clear that great care is advisable in applying the finite-difference techniques, it should be emphasized that the codes are quite capable of greatly advancing our qualitative and quantitative understanding of explosively generated stress waves in geologic media. In particular, I would like to discuss several types of studies that are basic to evaluating the current state of the art and establishing a credible prediction capability pertinent to the detection problem.

In discussions of the calculational capability, we always seem to be driven to an argument concerning code capabilities and the adequacy or inadequacy of material property test procedures and/or data. Those who obtain material property data are often alarmed by liberties taken by the theoreticians' attempt to force fit it into preconceived models for geologic media. On the other hand, the theoreticians generally cry out for more data for different states of stress, particularly unloading states of stress, that can be used to develop new material property models.

The above dialogue takes place between the theoretician and the experimentalist who obtains laboratory material property data. An even more heated argument sometimes takes place between the experimentalist and the theoretician who, using laboratory determined material property information, attempts to predict a specific field experiment. More often than not, pretest predictions based on code calculations have not been as successful as hoped, and post-shot calculations usually "explain" why in some plausible way. Once again the theoretical and experimental types lock horns--the calculators being supercritical of field-test data scatter (see Figure 19), and the field-experimental types being supercritical of the theoretician's warped view of an unreal world. In this matter, I am inclined to side with the field experimentalists and would toss similar darts at those who break small homogeneous samples of rock in the laboratory in the hopes of describing in situ rock properties. While current studies of intact rock specimens are producing a better understanding of basic constitutive relations for such samples, I am doubtful that they will directly impact the most pressing problem--namely, the dynamic response of in situ rock masses.

Personally, I believe that parametric studies investigating the sensitivity of computer material response to variations in the material model will identify the crucial material model parameters that must be determined by laboratory or in situ tests. Such studies may also suggest a significant redirection of effort in current material property testing. In addition to attempting to understand wave propagation phenomena in terms of classical constitutive models, we should also study basic mechanisms related to wave propagation in in situ rock. For example, attention should be focused on such features as the fact that the rock is jointed and the fact that rock (intact or in situ) cannot withstand tension. The following subsections will review several parametric studies indicative of the direction I believe we should follow in approaching the problem of interest.

Spherical Wave Propagation in Brittle Elastic Media

Consider the spherical wave propagation phenomena produced by the instantaneous loading of a spherical cavity within an elastic medium. Figure 34 shows typical radial and circumferential (hoop) stress spatial profiles for the case where the cavity is loaded by a step pressure and Poisson's ratio is 0.25.
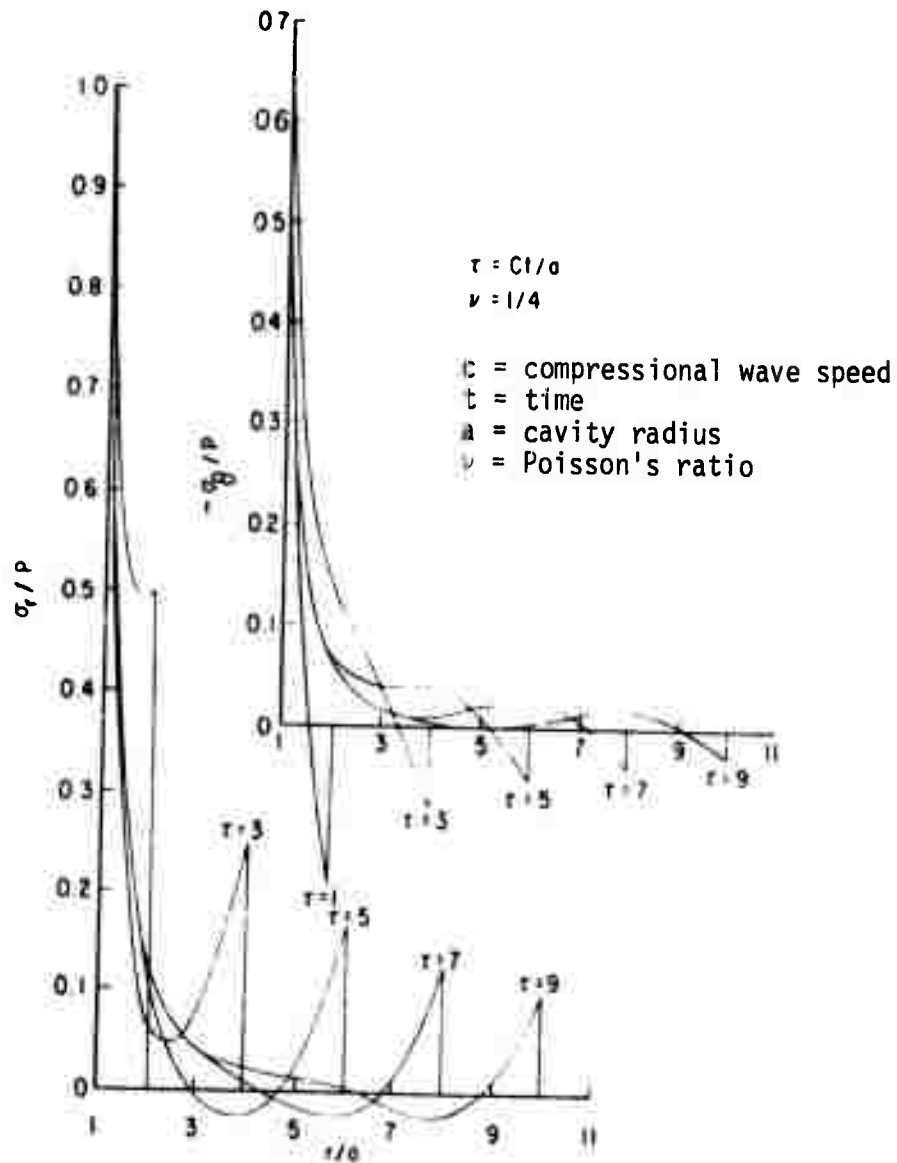
Figure 34.   Stress Spatial Profiles.

Note the tensile hoop stresses following the initial compressive phase--
particularly near the cavity at early times. These hoop stresses limit
the outward displacement much as the membrane stresses limit the outward
displacement of a spherical shell.

Since rocks fail under small tensile stresses, it is reason-
able to ask what dynamic behavior would follow if we simply allow the
medium to "crack" at the time hoop stresses exceed some small tensile

Figure 34. Stress Spatial Profiles.

Note the tensile hoop stresses following the initial compressive phase--
particularly near the cavity at early times. These hoop stresses limit
the outward displacement much as the membrane stresses limit the outward
displacement of a spherical shell.

Since rocks fail under small tensile stresses, it is reason-
able to ask what dynamic behavior would follow if we simply allow the
medium to "crack" at the time hoop stresses exceed some small tensile

value (σ_0). After the material cracks, the tensile strength is, of course, zero. We programmed the material model shown in Figure 35 into a spherically symmetric Lagrangian finite-difference code and proceeded to study the problem in a parametric way (26). Note that if the constitutive relations for the cracked and uncracked region were separately applied, the latter would lead to the usual elastic wave phenomena while the former would describe wave propagation down a rod of linearly varying diameter.

Although the primary motivation for this study was to study basic phenomena rather than to predict specific experimental results, we did choose elastic parameters and boundary conditions that would lead to an approximation of field measurements made on Piledriver, a 61-kt underground nuclear explosion in granite. Figure 35 describes the response that was generally observed. The motion field develops three distinct spatial regions, each separated by a cusp in the particle velocity (or stress) spatial profile. Immediately behind the front, the media is uncracked and the solution is identical to the homogeneous elastic solution. The second region is also uncracked, but the solution departs from the homogeneous elastic solution because signals propagate forward from the cracked region and alter the motion of the medium forward of the failure front. These signals are not initiated until cracking first occurs at the cavity wall. This fact, coupled with the fact that no wave can propagate faster than the compressional-wave speed, accounts for the first region being unaltered from the purely elastic solution. The third region is, of course, cracked, and its failure front expands at a rate less than the compressional wave speed of the uncracked media. Figures 36 and 37 compare particle velocity profiles from cases where the tensile strength is infinite (normal elastic) and essentially zero (total cracking).

Time histories at a given range are presented in Figure 38 for three values of the tensile strength and a fixed Poisson's ratio. Note the effect on displacement implied. This difference is better demonstrated in Figure 39, which gives the peak displacement attenuation as a function of the tensile strength. It should be noted that this variation in tensile strength causes no change in peak particle velocities.

As was pointed out previously, this study was motivated by a desire to understand spherical wave propagation in a material having brittle characteristics rather than a desire to precisely predict real material behavior. The model chosen had two parameters affecting the media response--the tensile strength and Poisson's ratio. Within the range of theoretically admissable values for these parameters, this very simple model can produce pulse shapes that are quite similar to those actually measured in field experiments.

It should be pointed out that, while the brittle failure model applied in these calculations is reasonable, the "elastic" behavior at the shock front is suspect at stress levels much higher than those considered here.

$$\text{FOR } \sigma_\theta > \sigma_0$$
$$\sigma_r = \lambda(\epsilon_r + 2\epsilon_\theta) + 2\mu\epsilon_r$$
$$\sigma_\theta = \lambda(\epsilon_r + 2\epsilon_\theta) + 2\mu\epsilon_\theta$$

$$\text{FOR } \sigma_\theta \le \sigma_0$$
$$\sigma_r = E\epsilon_r$$
$$\sigma_\theta = 0$$

$$\rho = 2.67 \text{ gm/cm}^3$$
$c = 18,000$ fps
 = compressional wave speed
t_f = time of initial cracking
v_f = velocity of failure front
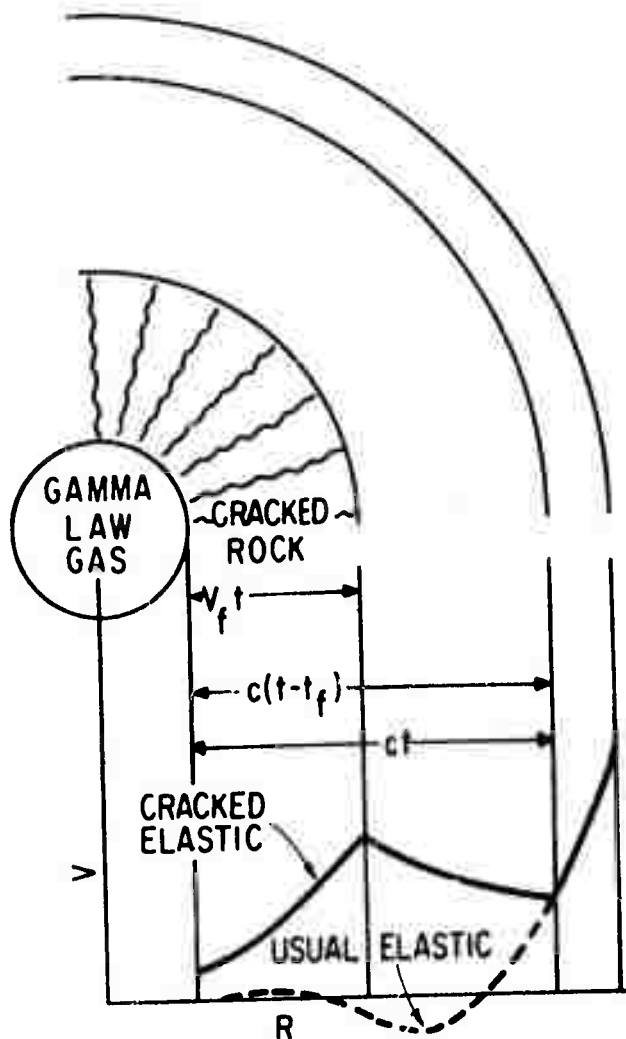
Figure 35. Brittle Elastic Rock Problem Description.
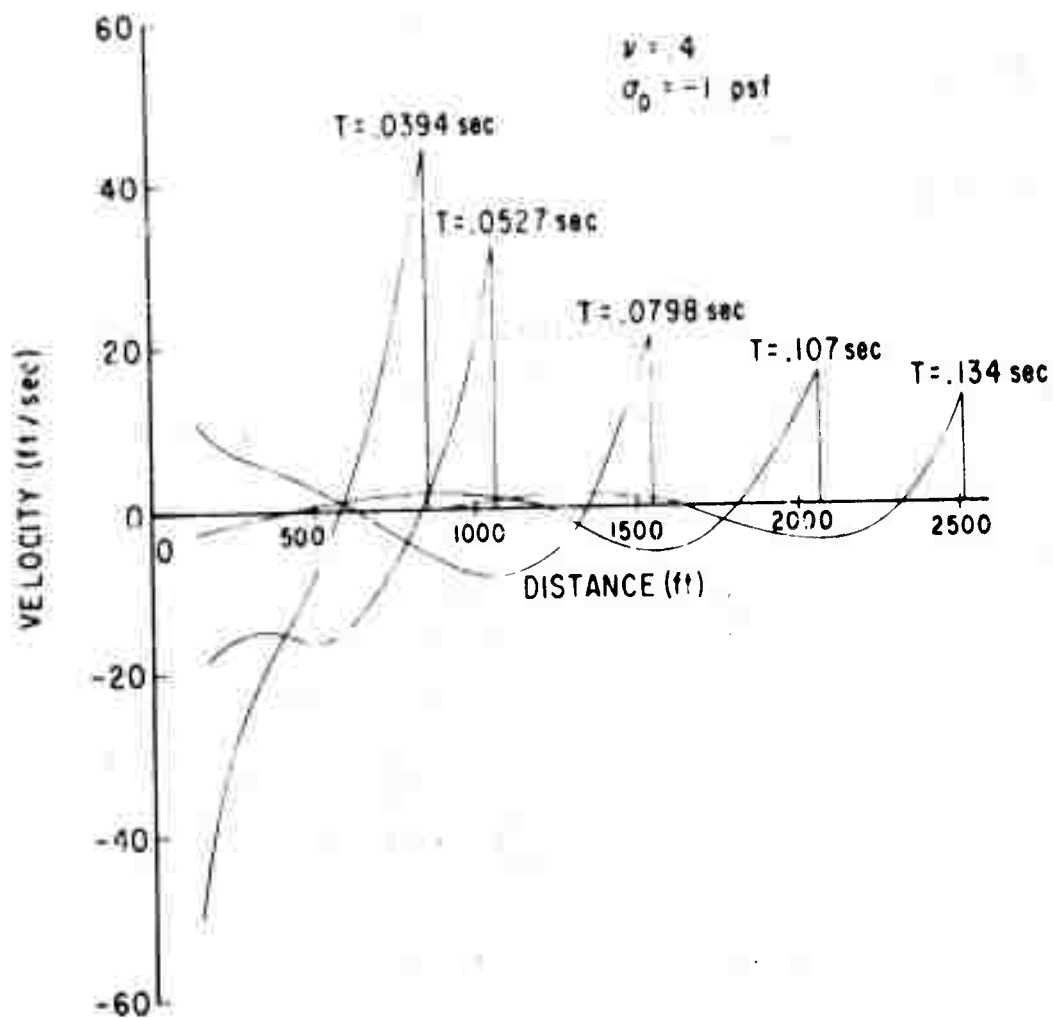
Figure 36.   Velocity Profiles--Usual Elastic Model.
Tensile strength = ∞.

90
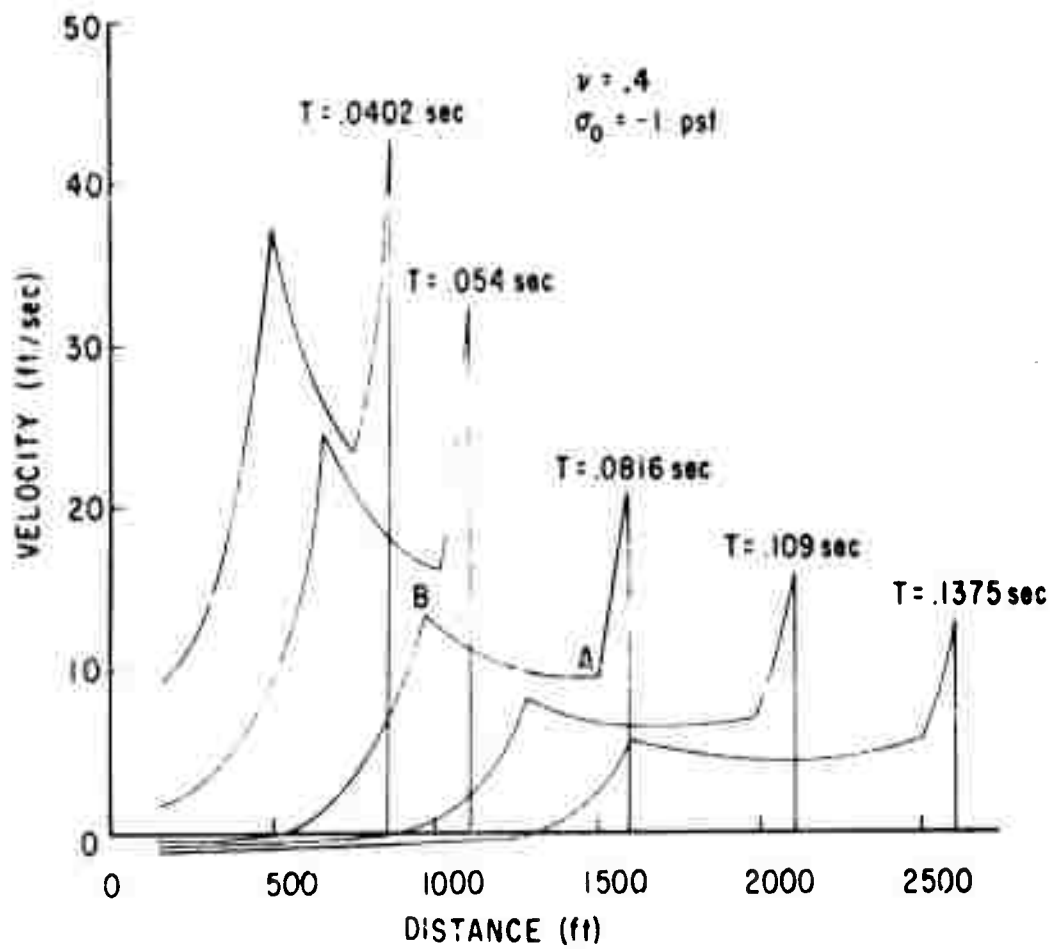
Figure 37.   Velocity Profiles--Total Cracking.   Tensile strength = 0.
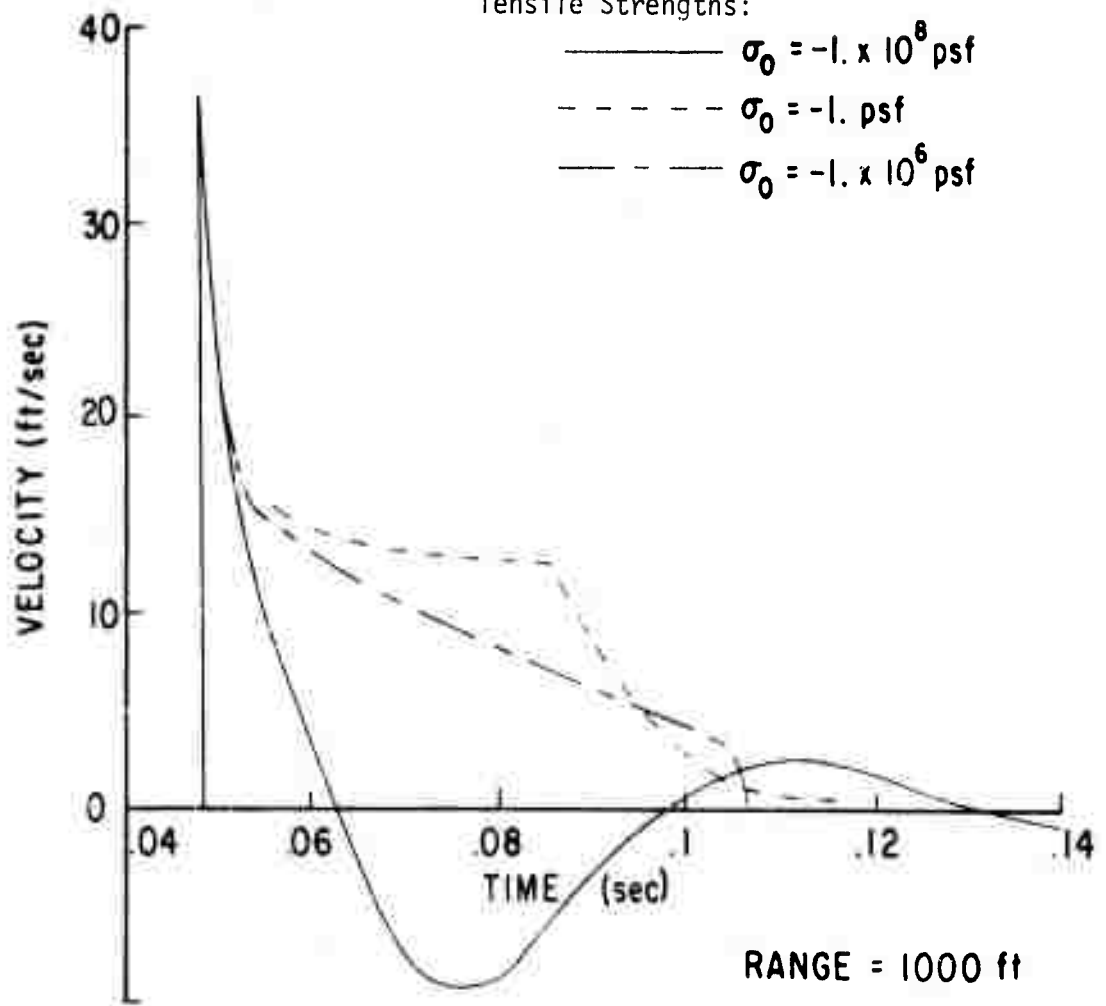
Figure 38. Particle Velocity-Time Histories at 1000-ft Range.
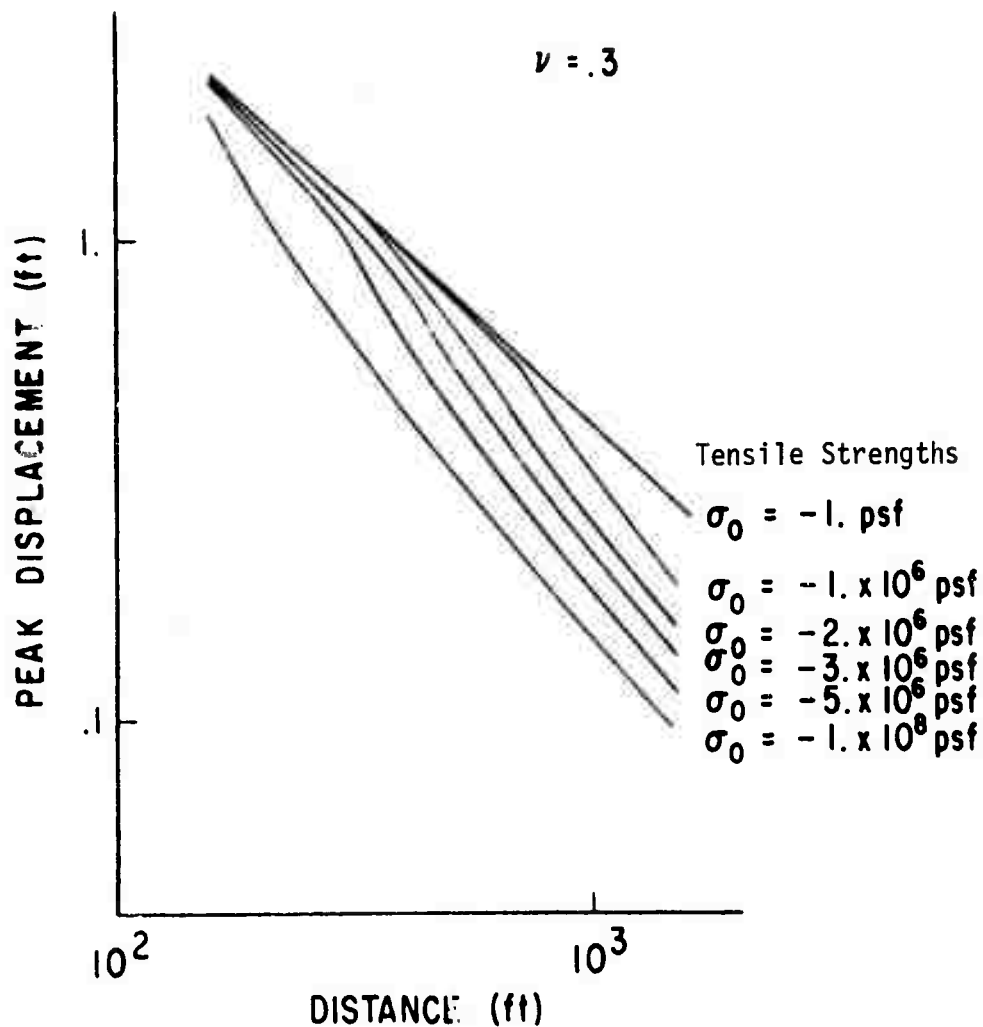
Figure 29. Particle Displacement Attenuation.

At sufficiently high stress levels, inelastic action surely occurs and more complicated models may be required. The following subsection will discuss two pertinent studies that use elastic-plastic models for the material behavior.

Spherical Waves in Inelastic Media

In a recent report, Isenberg, Bhaumik, and Wong (27) discuss several spherical wave calculations using baseline laboratory material property data for Cedar City tonalite. The models used in the calculations entailed a variable bulk modulus and a constant shear modulus. The parametric calculations performed included no-yielding,

93

yielding according to a von-Mises yield condition, yielding according to a Mohr-Coulomb yield surface and a plastic-potential flow rule, and yielding according to a Mohr-Coulomb yield surface and a Prandtl-Reuss flow rule. Instead of tensile failure criteria of the type indicated above, a restriction was placed on the mean principal stress such that 1900 psi tension could not be exceeded. It should be noted that such a condition can allow rather large tensile circumferential stresses. The various material model parameters are shown in Figure 40.

The calculations were performed with a pressure boundary condition to simulate several experiments performed by Physics International (28) involving the detonation of spherical chemical explosives in granite blocks. Figure 40 shows the effects of varying the yield parameters on the attenuation of peak radial stress. Figure 41 shows the circumferential stresses produced at various ranges in the four cases. With the exception of the Mohr-Coulomb/plastic potential case, unrealistic tensile stresses were developed. Figure 42 shows the sensitivity of the cavity wall displacements to the variations in yield criteria. Unfortunately particle velocity and displacements were not compared in general. The authors' conclusions of the features of their particular mathematical model that most strongly affected the stress results are:

(a) The amount of permanent volumetric compaction during hydrostatic loading and unloading;

(b) The amount of dilatancy accompanying inelastic deformation; and

(c) The yield criterion during unloading.

Laboratory and field tests could possibly be designed to address the specific areas pointed out as influential in stress-wave propagation phenomena, thereby suggesting which (if any) of the proposed material models is more influential.

The authors present a number of "stress trajectory" and stress-strain trajectory plots that trace the path followed in the various calculations. These plots suggest the stress and strain states most often undergone in the calculations and therefore identify areas deserving a concentration of material property testing effort.
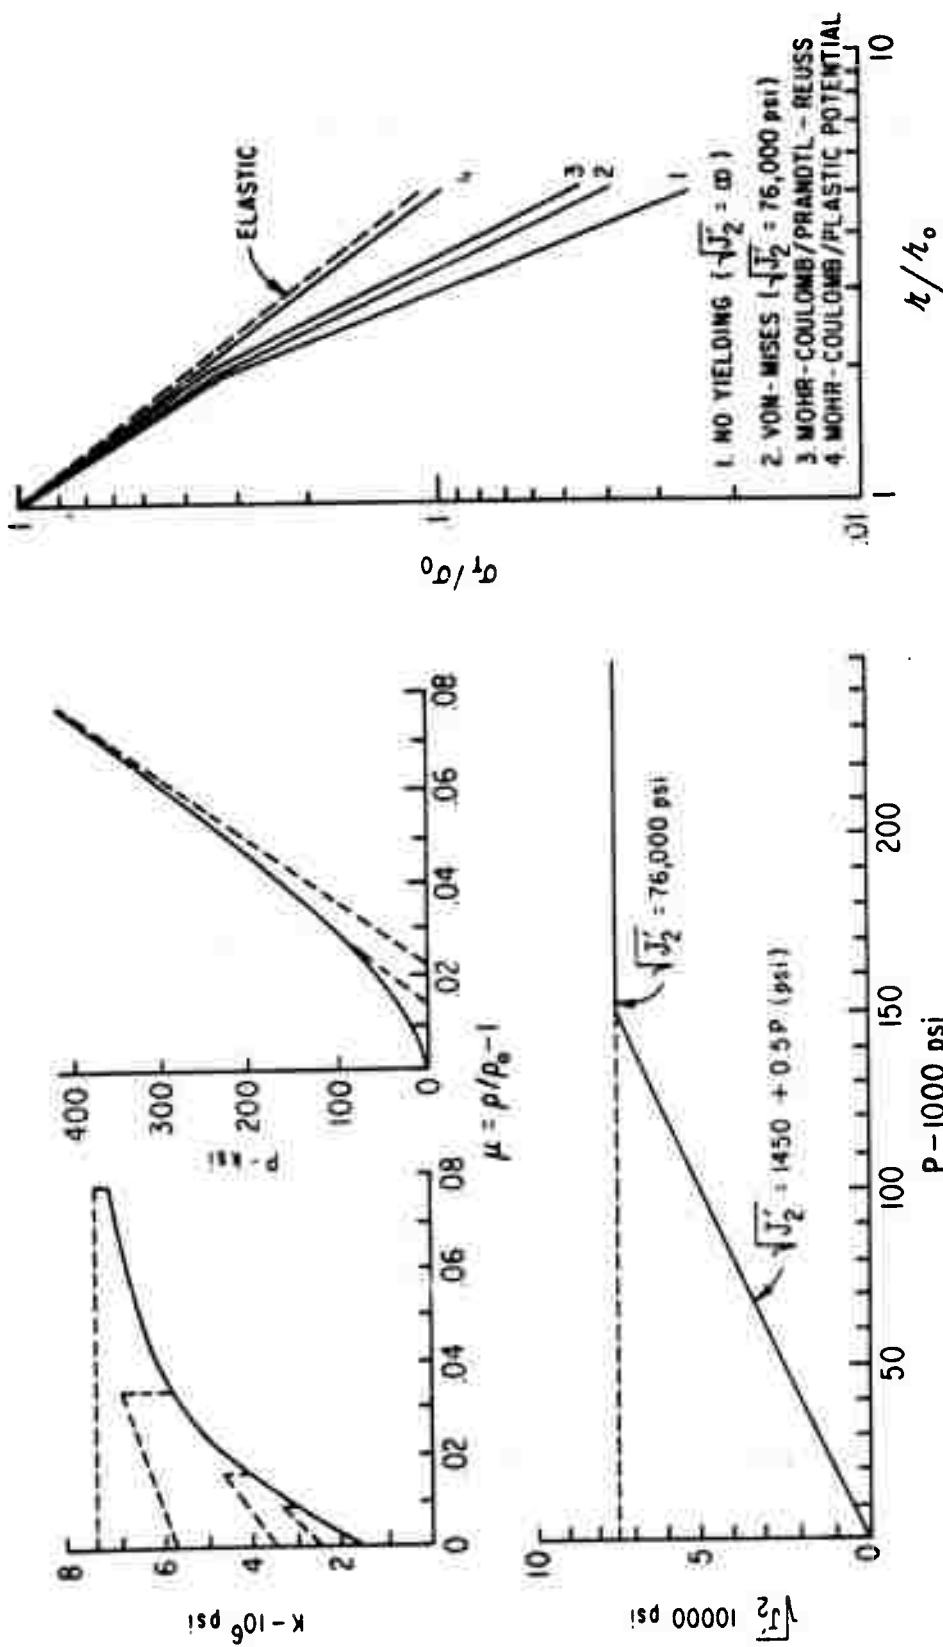
Figure 40. Effect of Yield Criteria on Peak Stress Attenuation. k = bulk modulus, σ_r = radial stress, σ_0 = initial stress at cavity wall.

95
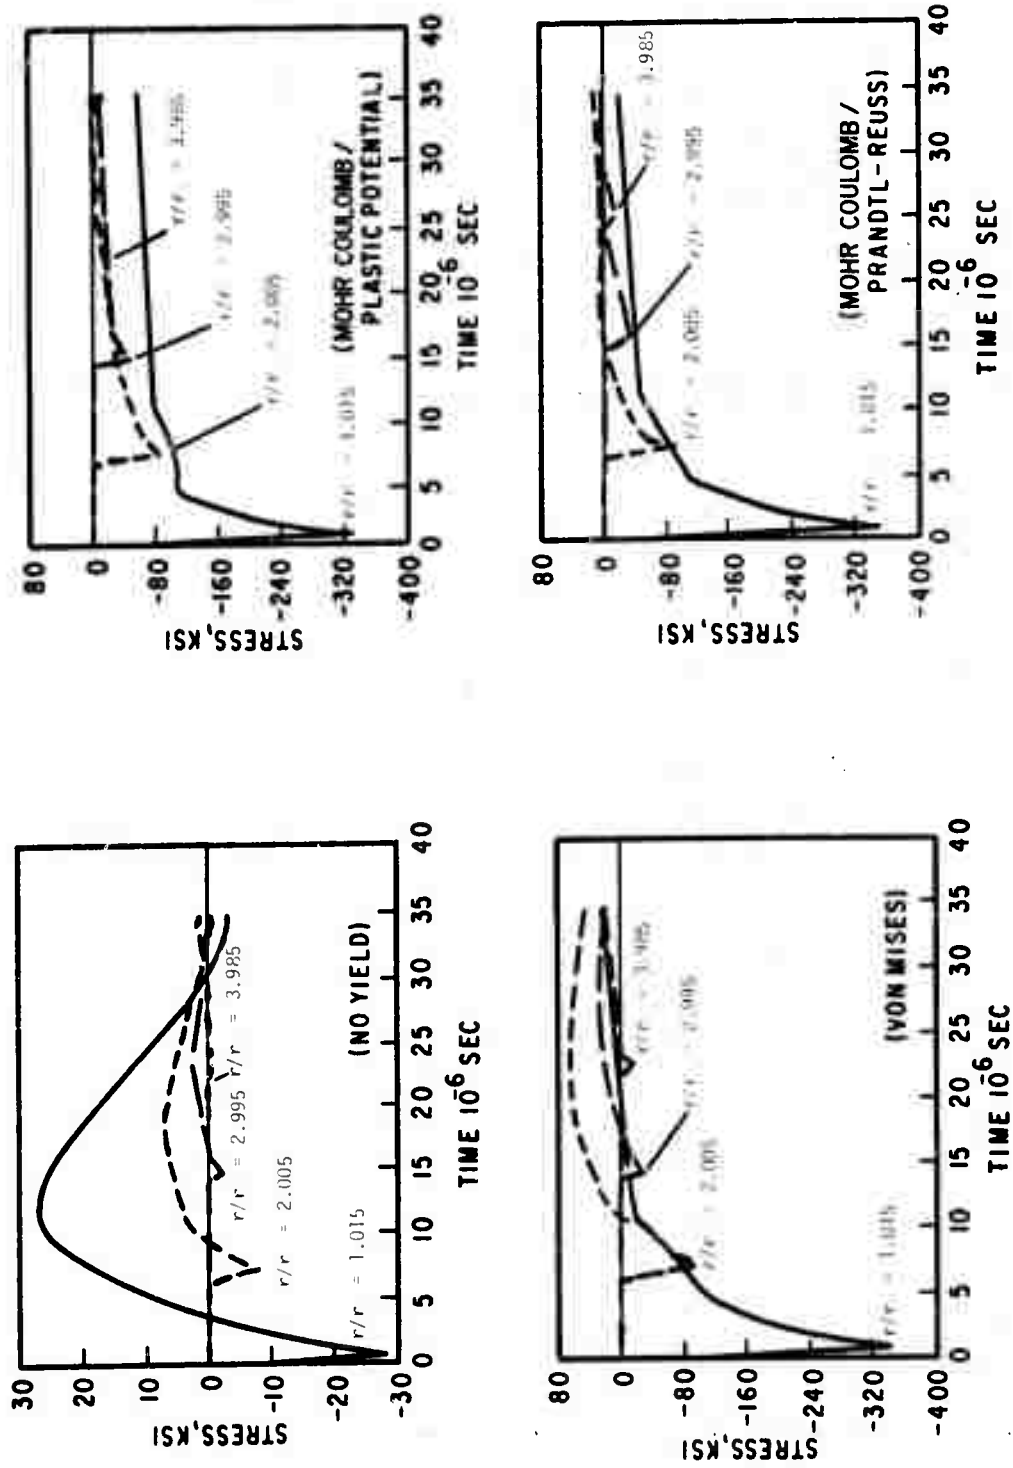
Figure 41. Circumferential Stresses.

96
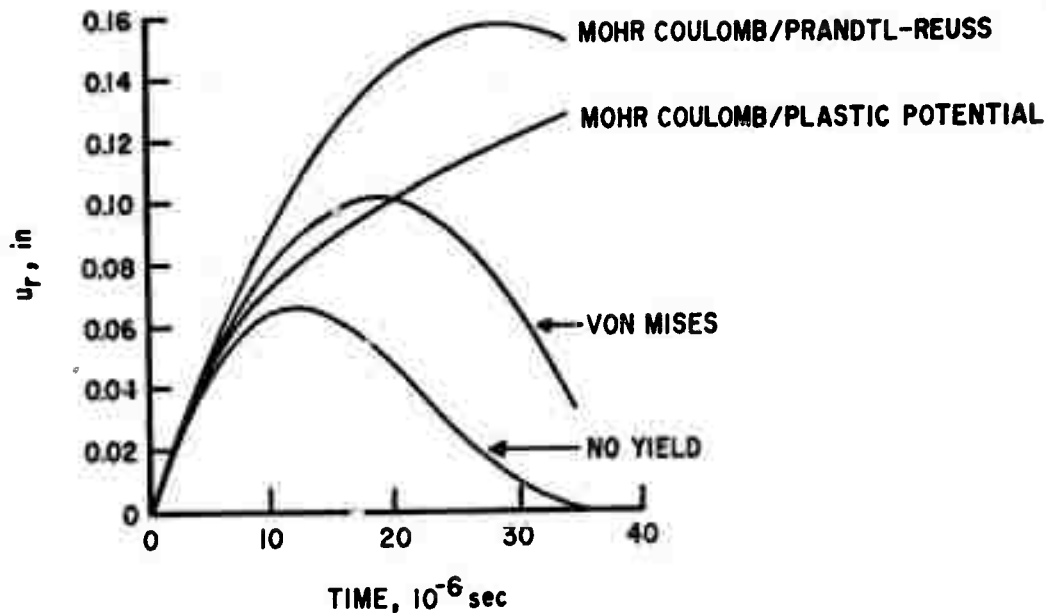
Figure 42. Displacement of Cavity Wall.

Another pertinent set of parametric results was presented in an interesting recent paper by Terhune, Stubbs, and Cherry (29). They indicated the variation in compressibility curves shown in Figure 43 as representative of variations they have observed at various Plowshare test sites. Note the variation indicated between a type A and type B rock. As in the previously discussed paper, particle displacements were not presented. Figure 44 indicates the expected variations in shear strength of various types of rock, and the drastic effect of such a variation on a particle velocity profile. Note that the cavity displacement is significantly affected by the variation between a "dry solid rock" and a "wet rock," while the difference between its displacement for a wet rock and a fluid (at the particular instant of time indicated) is rather trivial.

It should be noted that the yield surface used in these calculations is different than those discussed by Isenberg, et al. (27). As discussed by Cherry in a number of publications (see Ref 30, for example) this material model collapses extensional, compressional, and torsional failure data to a single failure surface expressed in terms of the third invariant of the stress deviator matrix in addition to the stress invariants usually included in the classical elastic-plastic formulation.
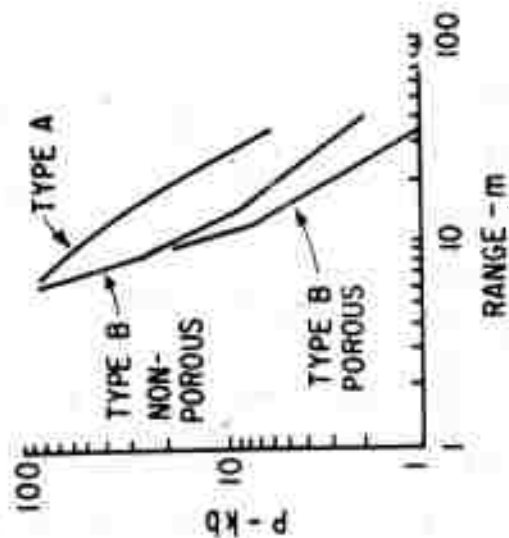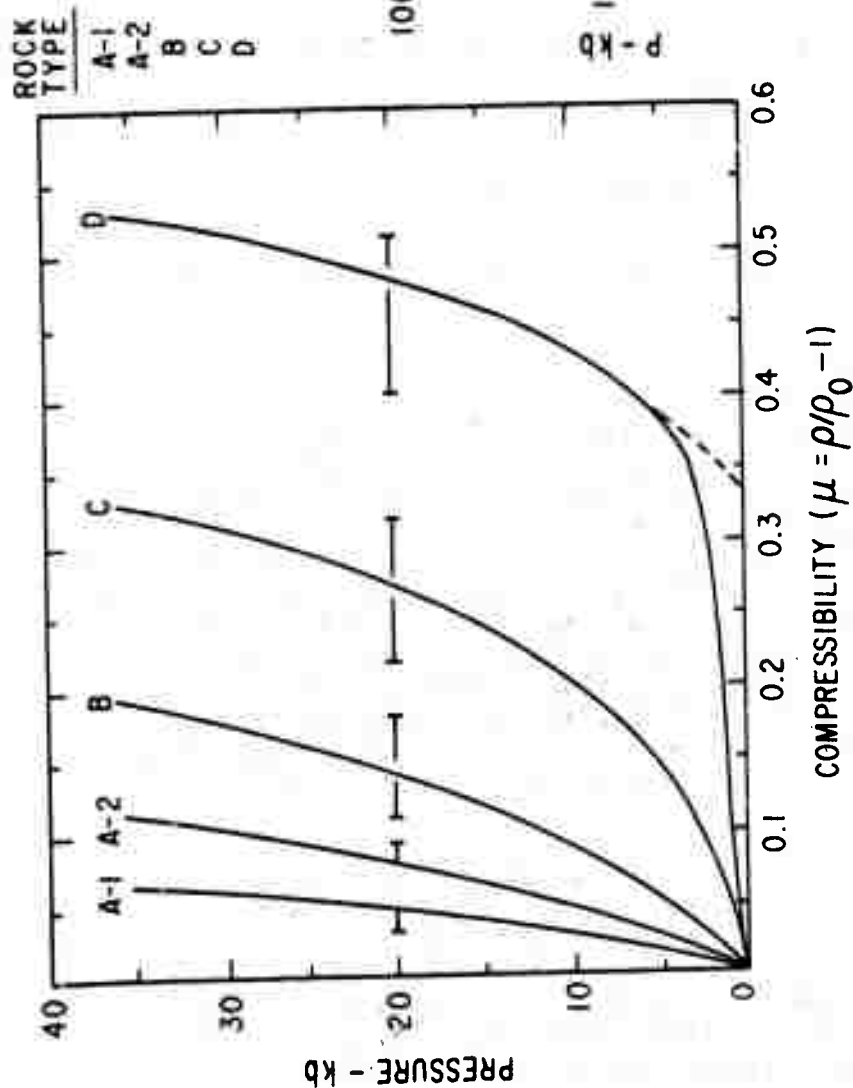
Figure 43. Compressibility Variations.
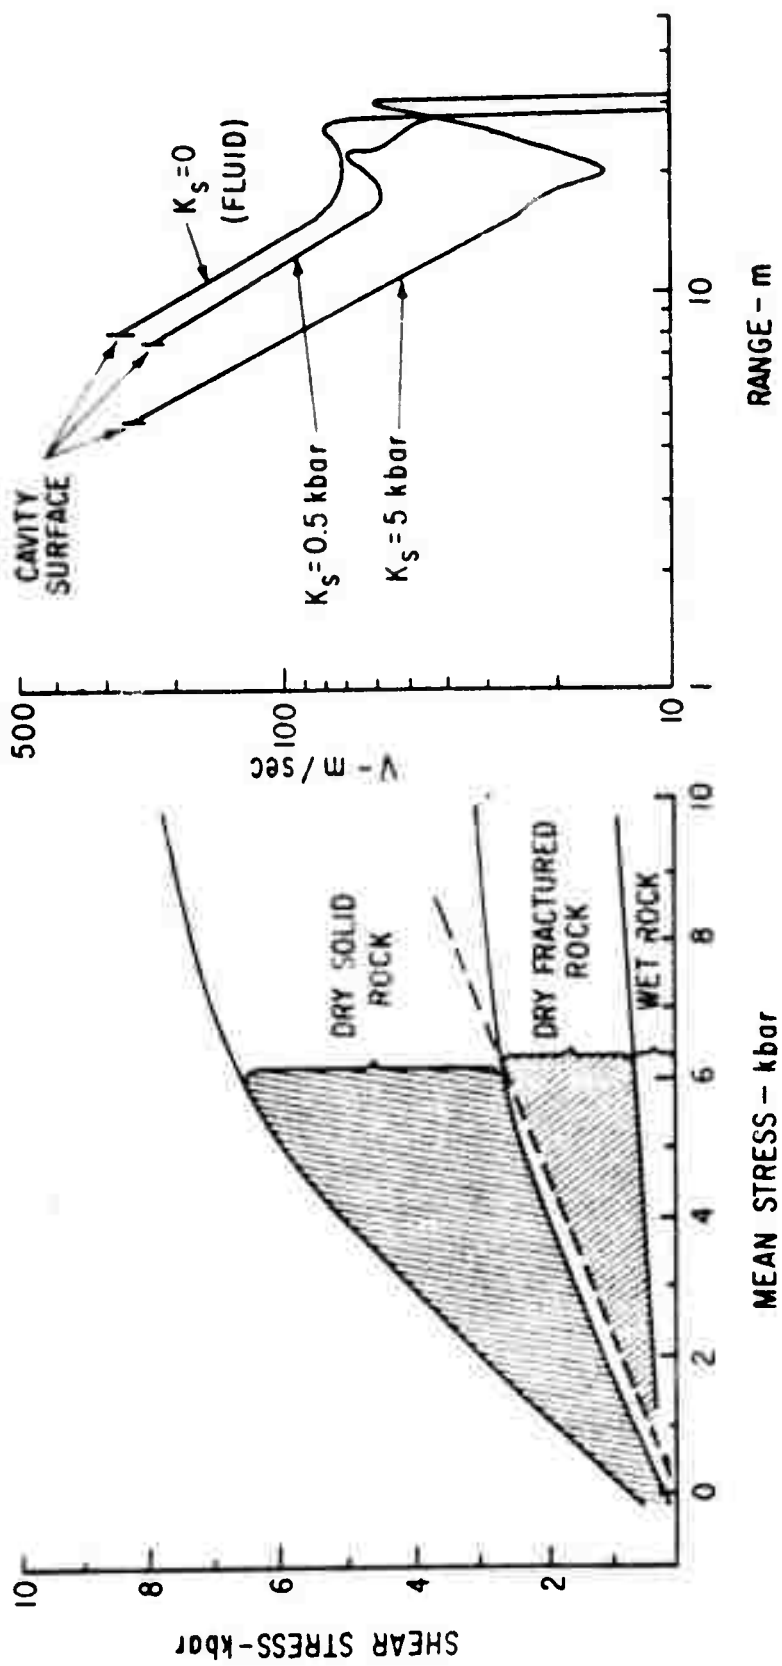
Figure 44. Shear Strength Variation

Effects of In situ Jointing

As indicated by Figure 20, the fact that rock is join _u can-not be ignored in the investigation of the dynamic response of in situ rock. Blocks of rock obviously can slide along pre-existing joint planes with a resulting displacement field more like that expected from pushing on a stack of bricks rather than what one normally expects from continuum motion. Actually nothing is really a continuum--whether or not one can treat a medium as being a continuum is really a question of scale. In my opinion, judging from what limited field data I have seen, the relative displacements indicated in Figure 20 are likely to be about the same as the nominal or average displacement occurring at the point in question. Noting that such motions can occur, we should not be at all surprised by the scatter in field data indicated in Figure 19. Conversely, attempting to compare theoretical results to one or two data points is rather absurd.

When one acknowledges the existence of such relative motions, then it is most pertinent to further question how one can make use of laboratory material property test data in describing such phenomena. Judging from the kind of motions indicated in Figure 20, it would seem reasonable to assume that the in situ shear strength is dominated by the resistance to sliding along jointing planes rather than by the strength of the intact rock. Compressional-wave phenomena is likely affected by diffraction phenomena and by the direct transmission of a shock across the joints via closing the pre-existing cracks.

As an initial step in addressing the relationship between in situ and intact rock, a theoretical study of the propagation of plane one-dimensional compressional waves through a cracked-rock model consisting of periodic transverse jointing has been conducted primarily by Abbott (31-34). The theoretical model for the study consisted of regularly spaced open cracks of width ΔL between elastic blocks of length L as indicated in Figure 45. A computer code (CRAC-1) was used to perform an extensive parametric study varying boundary conditions, crack widths, and joint spacing. Based on this study, a constitutive relationship was developed that represents the displacement response of a jointed medium in a continuum approximation. The resulting normalized stress-strain relation is given in Figure 46. Comparisons between CRAC-1 results and calculations performed with a Lagrangian code using the "continuum" model have been most encouraging. The model is now being evaluated by experiments in the laboratory.

Since the joints in rock are not random but have preferred directions, they will cause the in situ material to display an iso-tropic behavior both in its compressibility and its shearing strength. The eventual goal of this research is to be able to account for such behavior rather than to assume, as so often has been done in the past, that the field is made up of isotropic rock with the properties of intact specimens.

100
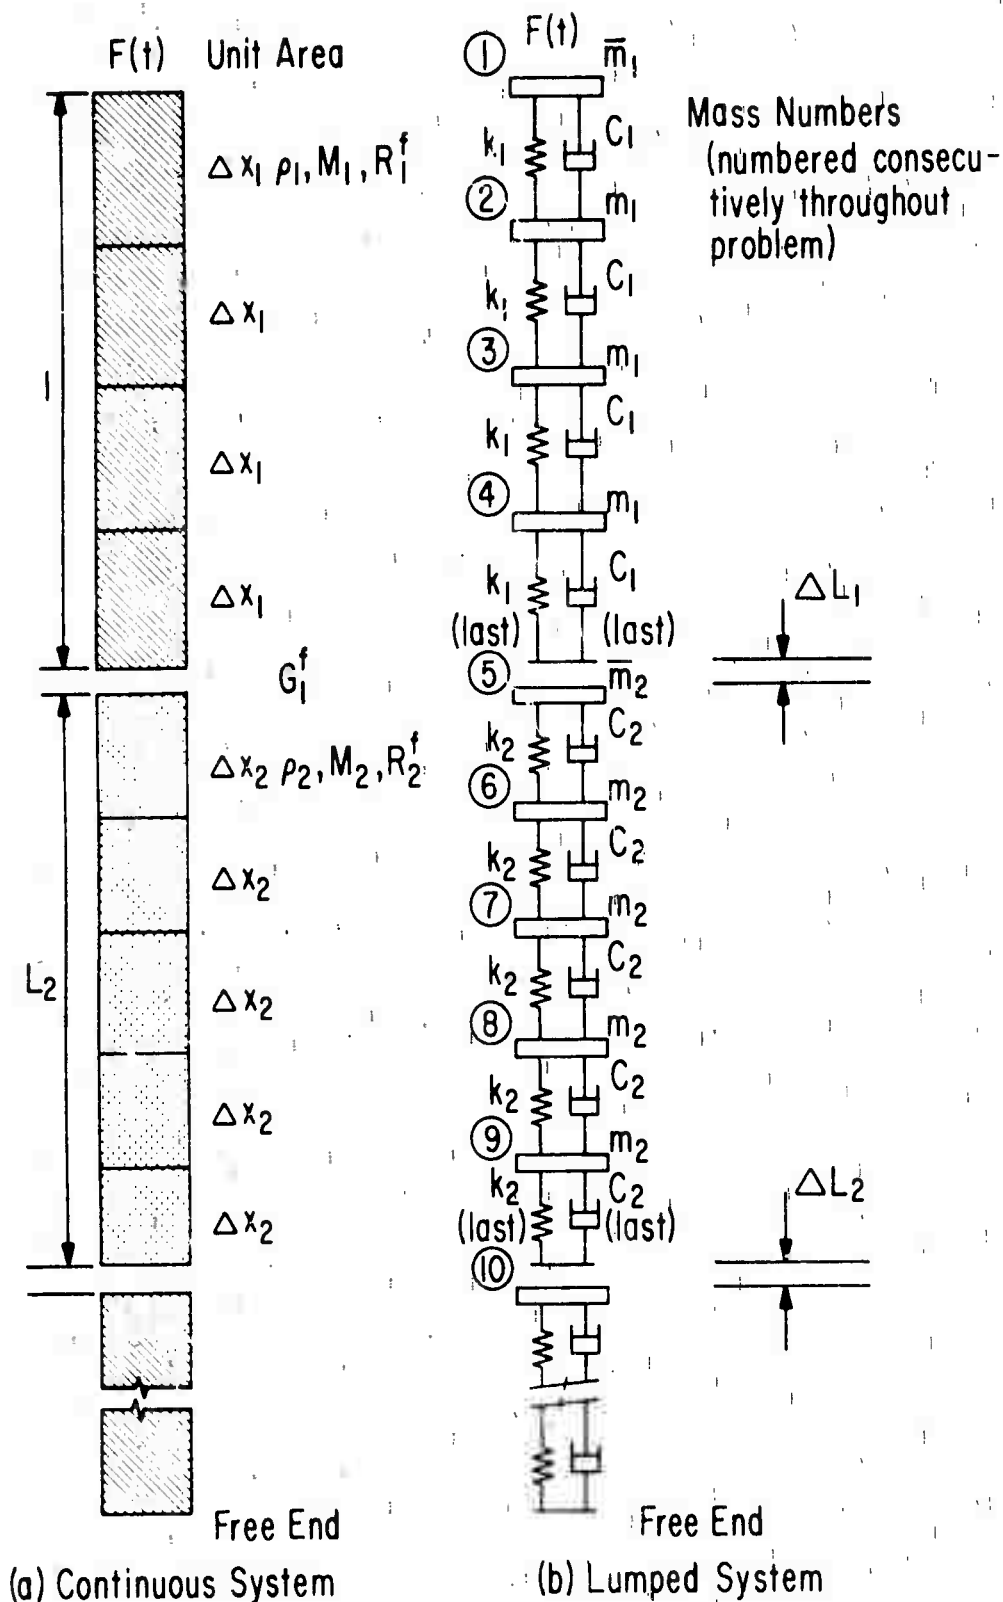
Figure 45. Lumped Parameter Model Used in CRAC-1 Program.
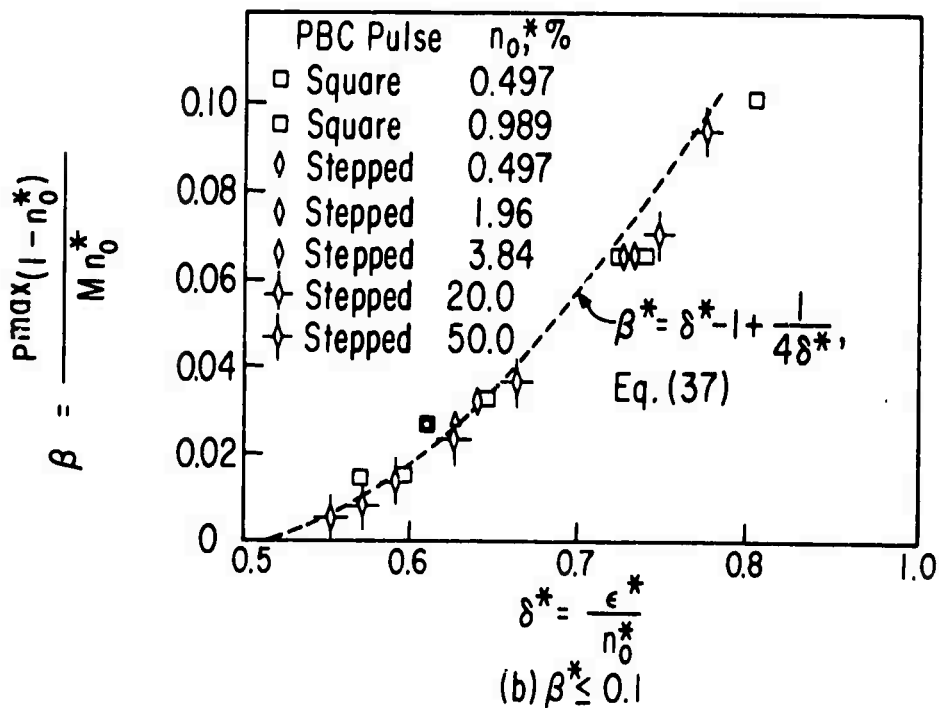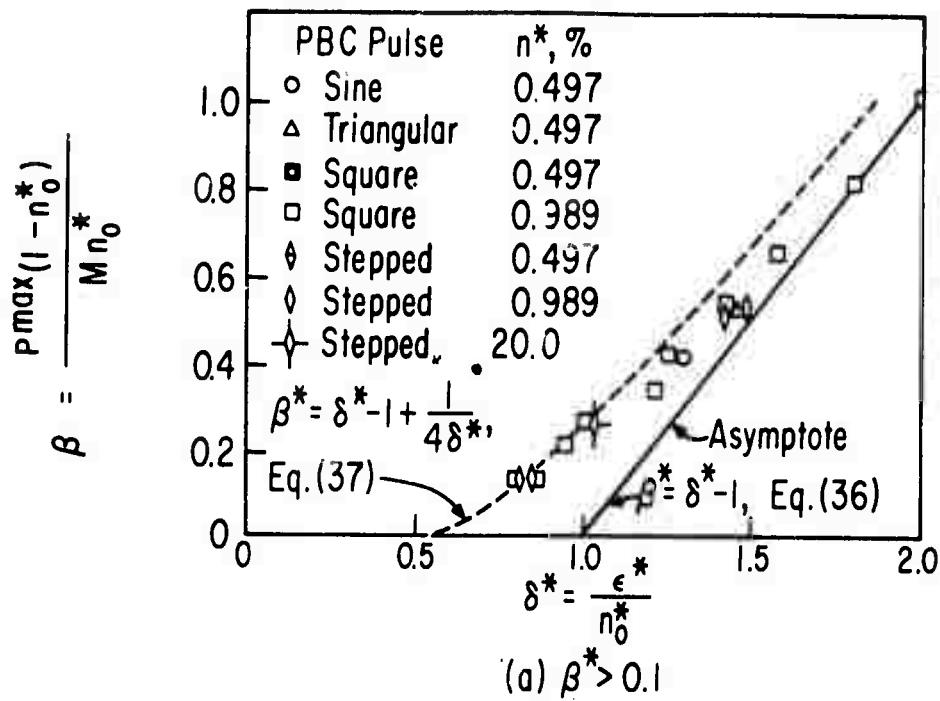
Figure 46. Normalized Stress-Strain Curve for Cracked Rock.
(Abbott, 33).

102

A vivid example of this anisotropy is shown in Figure 47. These photographs are the documentation of an interesting experiment performed one Sunday afternoon in an arroyo a few miles from Albuquerque. Several of us got to wondering what would happen if we were to set off an explosion in an orderly stack of bricks. Not having the financial support for such an experiment, and being a little impatient with the amount of red tape required to do the experiment anyway, we decided to invest in some 4032 sugar cubes and a few firecrackers. Figure 47 shows the result of this effort. The motion produced by the explosion shows a clear tendency toward taking the path of least resistance.

Recently, Abbott (35) has generalized the above mentioned model for cracked rock to a two-dimensional anisotropic constitutive relation for orthogonally cracked rock, e.g., as in the sugar-cube experiment. Figure 48 shows vector plots of the motion computed using this aniso-tropic material model. In comparing these results to the observations from the sugar-cube experiment, we are somewhat gratified by the quali-tative agreement.

In my opinion, studies such as these are pertinent applica-tions of the codes to better understand basic phenomenological mecha-nisms. They should be conducted in conjunction with controlled experiments and field tests to substantiate or reject implications of such theoretical models.

Summary

This section has indicated several parametric studies that either serve to evaluate the important parameters in existing material models or to address more basic questions that require answers in the formulation of more appropriate models for in situ rock. In my opinion, these are the most important kinds of applications that can be made with the existing calculational state of the art. I am convinced that only through such studies, interacting with controlled laboratory and field experiments will code calculations significantly improve our quantita-tive understanding of dynamic, in situ rock response to intense shock waves.

Conclusions

The following conclusions are based, in part, on the informa-tion referred to and contained in this paper. However, they represent my thinking as also influenced by other studies and experiments and even by the bias that I suppose I have developed over the past several years in this frustrating arena.

Figure 47. Explosion in a Stack of Sugar Cubes.
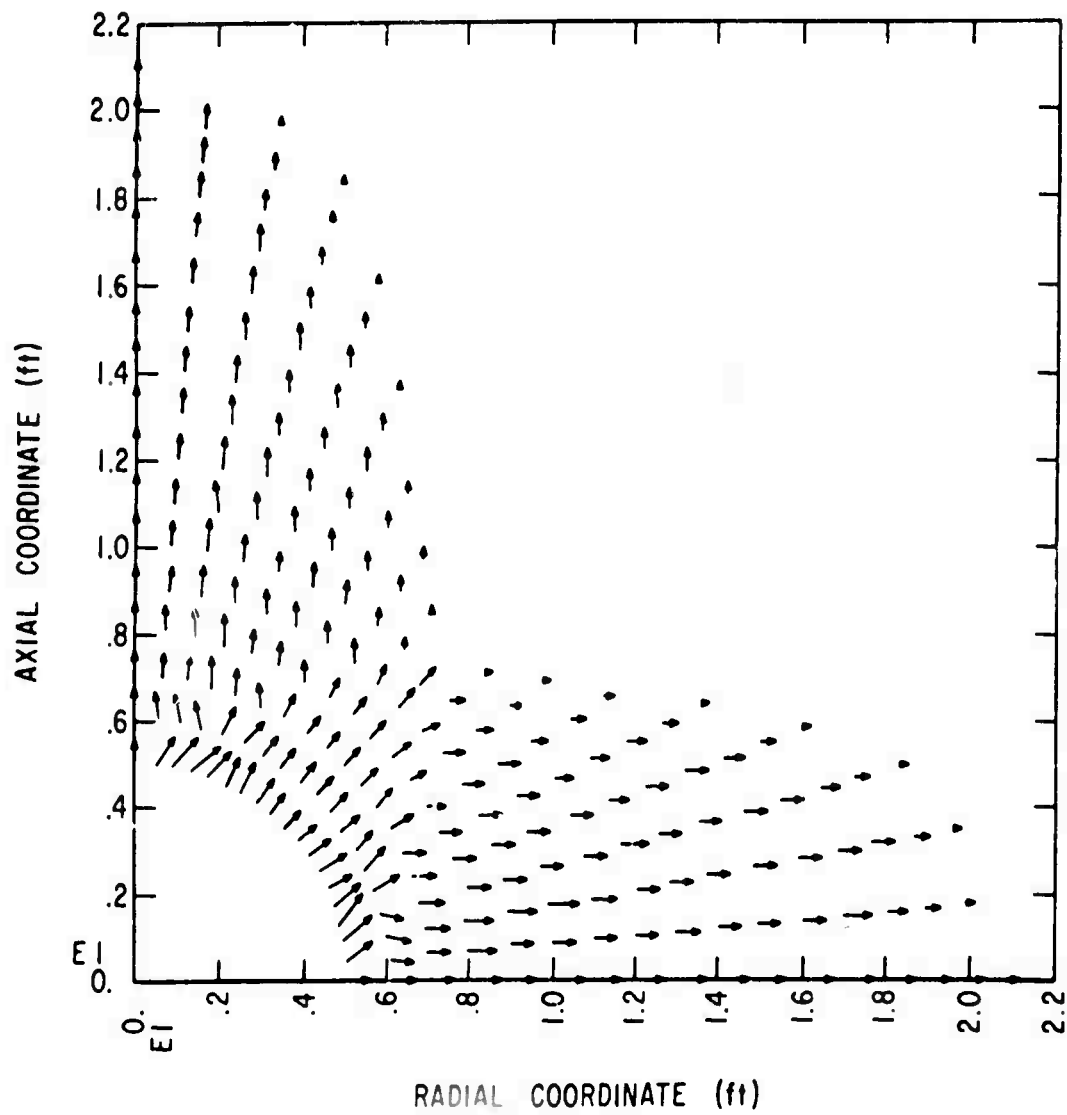
Figure 48. Vector Plots of Motion in Anisotropic Rock.

105

(a) Assuming that a systematic, quantitative, objective method for predicting ground shock produced by nuclear explosions is desired, then several prevalent attitudes about the current state of the art must change. First, it must be understood that numerical errors can be significant, particularly for two-dimensional calculations involving many orders of magnitude in stress-wave attenuation. In line with this comment, it should be understood that accurate calculations are likely to be expensive. Second, the vocabulary associated with the computational effort must change to involve more quantitative, scientific terminology and fewer qualitative generalities. For example, there should be some objective demonstration that the zoning is "adequate" for any calculation that is purported to be quantitative.

(b) While numerical errors can cause serious problems, careful application of one-dimensional spherical codes should produce results sufficiently accurate to solve at least the early stages of the motions produced by underground explosions in homogeneous media. If inhomogeneities (discrete layering, smoothly varying properties, or local conditions) are important in influencing the late-time motion, then two-dimensional effects must be considered in the calculations. In my opinion, existing two-dimensional codes (if carefully applied) should be sufficiently accurate to provide qualitative understanding and, in some cases, quantitative results. A key variable in achieving accuracy is generally the requirement for a large number of zones, which, in turn, raises the cost of a given calculation.

(c) Perhaps the greatest folly of all of the work that has gone into code development and the calculational effort has been the short-sighted search for numbers in the form of predictions rather than systematic studies that are aimed at gaining insight and the understanding of basic phenomena. In my opinion we have in the main wasted much time and money force-fitting preconceived theoretical models for geologic media to laboratory data from a relatively small number of applicable stress states rather than attempting to understand the basic mechanisms that control the late-time behavior on in situ rock.

(d) Parametric studies such as those indicated in this paper are most useful in identifying the key material property parameters that control late-time motions. In particular these and other studies have indicated the importance of: (1) compressibility, (2) porosity, (3) water content, (4) strength, or yield criteria including the flow rule (assuming an elastic-plastic model).

(e) The effects of in situ joints on explosively produced ground motions have not received enough attention. Experimental work has clearly demonstrated that late-time phenomena is very much influenced by pre-existing joint patterns. It is noted that the preferential directions of such jointing leads to a condition of anisotropy for in situ rock.

106

Recommendations

Recognizing that ground shock close in to completely contained bursts is controlled by the mechanical and thermodynamic properties of in situ geologic materials, acknowledging that in situ material properties are likely to be somewhat inconsistent with laboratory determined material properties, and assuming that only qualitative geologic and chemical information will be available for foreign sites where clandestine testing is conducted, I believe that the careful application of state-of-the-art finite-difference techniques can provide the "cornerstone" for a systematic theoretical and experimental program to significantly improve confidence in our prediction capability. Such a recommended program must include basic material property (especially in situ) testing, basic phenomenology studies, and confidence tests.

Phenomenology Studies

The driving force in such a program must be the basic phenomenology studies whose objectives are to: (1) identify material property and geologic structure parameters that control close-in, late-time phenomena, (2) determine the range of effects on ground motions implied by the uncertainties in specifying these parameters (based on existing information), (3) provide the technical guidance required to design meaningful laboratory and field experiments to provide more definitive estimates of the in situ material properties where the existing uncertainties are unacceptably large, and (4) correlate and evaluate the results of such experiments.

The first task in this study area would be to quantify the implications of currently existing material models through extensive parametric studies. Of primary importance in this task is the ultimate goal of reducing the number of parameters required to formulate adequate theoretical models to those that can reasonably be estimated from the limited available information for foreign sites. In particular, it is recommended that parametric studies be performed to: (1) determine the requirements for detailed Hugoniot data, including release adiabats, (2) determine the relative importance of compressibility, porosity, and strength (as defined by all current models for brittle or ductile material failure—including elastic-plastic flow rules), and (3) study and better define our basic understanding of basic mechanistic effects (jointing, etc) including the late-time effects of anisotropy.

Controlled laboratory experiments can also provide a better phenomenological understanding of late-time, in situ rock response. Systematic dynamic experiments on model rock may provide a much improved understanding of in situ jointing effects. Studies in "homogeneous" brittle materials may also be informative. Although I personally feel that in situ material property testing should have clear priority, a moderate laboratory material-property experimental program, especially tests on cracked samples, should be continued in parallel with greater emphasis on in situ testing.

107

Material Property Testing

As indicated above, I believe that the development of meaningful in situ tests to determine in situ material properties should receive a high priority. Since the late-time ground motions are probably most influenced by the jointing, such in situ tests should involve a reasonably large test bed area to include a significant number of joints.

It should be noted that previous high explosive and nuclear ground shock data from underground tests are themselves an in situ test of sorts. Therefore, it is recommended that a comprehensive compendium of such test data be prepared. The data presentation should include time-history details as well as peak value tabulations. Also, all material property information for the media in which the reported experiments were performed should be included.

As results from the phenomenological studies raise questions concerning in situ property details, field experiments should be conceived and fielded to resolve important issues. For example, at present I believe there are sufficient questions involving details of the yield surface, flow rules, dilatency, and brittle characteristics implied by existing theoretical models to justify a field experiment to resolve the question if such an experiment can be conceived. It is important to note that the theory should drive such experiments--and pretest predictions are of the utmost importance in establishing a credible prediction procedure.

Confidence Tests

The success of any program that purports to predict late-time ground shock is the development of confidence in the theoretical procedures and in the adequacy of the material property models used to represent in situ geologic media. As suggested above, the results of existing and possible future ground motion experiments serve as an indicator of the in situ material properties of importance. Existing computer codes dictate the mathematical tools and, to some extent, the mathematical models with which the behavior of real rock is likely to be described in the near future. Confidence in any prediction technique strongly hinges on just how well theoretical calculations match experimental results. Thus, the various calculational techniques should be correlated with the compendium of experimental results suggested above. If calculations match experimental results within the tolerances set by overall numerical and experimental error estimates, then these error estimates serve to define confidence factors. If they do not, then the cause of the discrepancies must be identified and fed back into the phenomenology studies.

In compiling the compendium of relevant test data, it should become obvious if significant gaps exist, and a field experiment may be

108

suggested to supply the missing information. Such an occurrence is a major confidence test, and published pretest predictions from all key contractor and government theoreticians should be mandatory.

Because ARPA's real problem involves the case where little geologic information and practically no material property data will be available, basic phenomenological information must be obtained to allow an extrapolation from the existing technology base to address these cases. Therefore, at some appropriate time, I suggest the following confidence test. Provide the calculators with the sketchy baseline information that might be available from Russian sites (density, depth to bed rock, chemical make up, etc) and task them with making predictions on that basis; then do an experiment and see how well we do in the pretest predictions.

A more modern comparison study than that presented on pages 78-85 might be useful in evaluating the current state of the art. Such a study should include strength characteristics and geologic layering.

References Cited

1. Trulio, J. G., Unpublished Presentation of Piledriver Calculations, 1968 DASA Long Range Planning Meeting, Jan 1968; Also ATI/AJA Technical Progress Report No. 3 to SAMSO, Beneficial Site Models Ground Motion Calculations, Contract F004694-67-C-0120, 18 Dec 1967.

2. Godfrey, C. S., Calculations and Experiments on the Dynamic Properties of Rocks, *Proceedings of the DASA Strategic Structures Vulnerability/Hardening Long Range Planning Meeting, DASA 2128*, 14-16 Jan 1969.

3. Wagner, M., and N. A. Louie, Hard Hat/Piledriver Ground Motion Calculations, *SAMSO-TR-69-47*, March 1969.

4. Gauvenet, A., Experiments with Underground Nuclear Explosions in the Hoggar Massif in the Sahara, *Symposium on Engineering with Nuclear Explosives*, Las Vegas, 14-16 Jan 1970.

5. Chabai, A. J., Synthesis of Shock Hugoniots for Rock Materials, Proceedings, Fifth Symposium on Rock Mechanics, Pergamon Press, New York, p. 347-366, 1963.

6. McQueen, R., Shock-Wave Data and Equations of State, in Seismic Coupling, edited by G. Simmons, Proceedings of ARPA sponsored meeting held at SRI, June 1968.

7. Gogolev, V. M., V. G. Myrkin, and G. I. Yablokov, Approximate Equation of State of Solid Bodies, *J. of Appl. Mech. and Tech. Physics*, No. 5, 1963 (Translation *FTD-MT-64-61*, Feb 65).

8. Trulio, J., and K. Trigger, Numerical Solution of the One-Dimensional Lagrangian Hydrodynamic Equations, *UCRL-6267*, 1961.

9. Trulio, J. G., and K. R. Trigger, Numerical Solution of the One-Dimensional Hydrodynamic Equations in an Arbitrary Time-Dependent Coordinate System, *UCRL-6522*, July 19, 1961.

10. Trulio, J. G., Studies of Finite Difference Techniques for Continuum Mechanics, *WL-TDR-64-72*, December 1964.

11. Trulio, J. G., Theory and Structure of the AFTON Codes, *AFWL-TR-66-19*, Jun 1966.

12. von Neuman, J., and R. D. Richtmeyer, A Method of the Numerical Calculation of Hydrodynamic Shocks, *J. Appl. Phys.*, v. 21, March 1950.

13. Amsden, A. A., The Particle-in-Cell Method for the Calculation of the Dynamics of Compressible Fluids, *LA-3466*, June 1966.

14. Private Communication from Captain E. A. Nawrocki of the Air Force Weapons Laboratory.

15. Papetti, R. A., Private Communication, May 1970.

16. Cooper, H. F., Comparison Studies of Finite Difference Results for Explosions on the Surface of the Ground, *AFWL-TR-67-25*, May 1967.

17. Johnson, W. E., OIL, A Continuous Two-Dimensional Eulerian Hydrodynamic Code, General Atomic Report *GAMD-SS80*, Jan 7, 1965.

18. Harlow, F. H., Two-Dimensional Hydrodynamic Calculations, *LA-2301*, Sept 1 59.

19. Andrews, D. J., C. S. Godfrey, E. Teatum, and E. T. Trigg, Calculations of Underground and Surface Explosions, *AFWL-TR-65-211*, June 1966.

20. Tillotson, J. H., Metallic Equations of State for Hypervelocity Impact, General Atomic Report *GA-3216*, July 18, 1962.

21. Trulio, J. G., J. J. Germroth, W. J. Niles, and W. E. Carr, Study of Numerical Solution Errors in One- and Two-Dimensional Finite Difference Calculations of Ground Motion, *AFWL-TR-67-27*, Vol I, 1967.

22. Sonnenburg, P., Private Communication, May 1970.

23. Cooper, H. F., Generation of an Elastic Wave by Quasi-Static Isentropic Expansion of a Gas in A Spherical Cavity; Comparison between Finite Difference Predictions and the Exact Solution, *AFWL-TR-66-83*, Sept 1966.

24. Godfrey, C. S., D. J. Andrews, and E. V. Trigg, Prediction Calculations for Free Field Ground Motion, *WL-TDR-64-27*, May 1964.

25. Trulio, J. G., W. E. Carr, and J. J. Germroth, Optimum Coordinate Systems for One-Dimensional Finite Difference Calculations, *AFWL-TR-67-27*, Vol II (To be published).

26. Burford, L., H. F. Cooper, and J. C. Thompson, Some Preliminary Calculations of Spherical Wave Propagation in Brittle Materials, AFWL Technical Memorandum *WLC-TM-70-004*, Jan 1970.

27. Isenberg, J., A. K. Bhaumik. and F. S. Wong, Spherical Waves in Inelastic Materials, *DASA 2404*, March 1970.

28. McKay, M. W., and C. S. Godfrey, Study of Spherically Diverging Waves in Earth Media, *DASA 2223*, March 1969.

111

29. Terhune, R. W., T. F. Stubbs, and J. T. Cherry, Nuclear Cratering on Digital Computer, Presented at Proceedings of the ANS Topical Meeting, Engineering with Nuclear Explosives, Las Vegas, 14-16 Jan 1970; also *UCRL-72032*, Feb 1970.

30. Cherry, J. T., D. B. Larson, and E. G. Rapp, A Unique Description of the Failure of a Brittle Material, *UCRL-70617*, Sept 1967.

31. Abbott, P. A., and R. O. Davis, One-Dimensional Wave Propagation through Cracked Rock, A Description of the Crac-1 Program and Problem Solutions, *AFWL-TR-69-16*, May 1969.

32. Cooper, H. F., P. A. Abbott, and N. R. Mitchell, Fundamental Studies in One-Dimensional Wave Propagation; Proceedings, Strategic Structures Research Vulnerability/Hardening Long-Range Planning Meeting, Vol I, *DASA 2288-1*, June 1969.

33. Abbott, P. A., Dynamic One-Dimensional Continuum Model for Transversely Cracked Rock, *AFWL-TR-70-3*, March 1970.

34. Abbott, P. A., and R. O. Davis, Development of a Dynamic Continuum Description for Cracked Rock, Submitted for presentation at the 12th Symposium on Rock Mechanics, 1970.

35. Abbott, P. A., Dynamic Two-Dimensional Continuum Model for Orthogonally Cracked Rock, AFWL-TR- to be published.

DISCUSSION OF CODE CALCULATIONS

CHAIRMAN SIMMONS: Do we have comments from the panel?

MR. TRULIO: To get the ultimate in credibility, when present test calculations are made they ought to be locked up for a while until the measurements are all in.

MR. COOPER: I think they should be reported pretest, too.

I had an interesting exercise. We were involved with a program recently of actually conducting experiments where we would sit down and try to predict what was going to happen in that experiment. I have played with computers and am very interested in that, but when it came time to actually predict what was going to happen, I didn't rely on computers and I didn't rely on code. It was an interesting retrospection, if you will, because I didn't believe the code would give me the answers I wanted, particularly the displacement. The displacement are late-time phenomena and they are inherently different to predict than the other.

MR. TRULIO: I would like to make some general comments. First, the accuracy of a numerical method--even the slides that you show suggest to me, having looked at shock data and pretest predictions compared to it, that while there aren't negligible errors in the calculations, the uncertainties in the mechanical property representations from geologic media are still much bigger sources of difficulty. So far as studies on accuracy of numerical methods themselves are concerned, I think there are a few cases that you can do analytically. I am sure they haven't all been done. One of the most important since you reported the results of your calculations would be an exact solution for the one-dimensional linear case for a half space. The only case that I know of having been solved is wave propagation in an infinite, one-dimensional medium. I think the problem to be soluble for a half space. You could find out what numerical boundary really is the best. It is not at all clear that that would be the exact solution.

There is some information on the rate of convergence of these methods and it might not be necessary, if you want to make an error-bound estimate, to solve a whole series of problems. The general rate of convergence with mesh point density I think is pretty well established and discouragingly slow--but it is pretty well established.

The way it could be used is to simply perform two different meshes that would permit you to extrapolate or extract from the changes in the answer going from one mesh to another and get a rough estimate of the error in both, and of course then a rough estimate of the exact ends.

The model of the mediums, I think, is still by far the biggest source of uncertainty in the calculated results. So far, I would say what has happened is we have done a zero order of approximation on a continuum model for a medium that everybody knows isn't--and also a nonhomogeneous medium, but isotropic certainly, that has cracks in the medium. The right step is one of considering the cracks to be isotropic, distributed in some random way such that there is no preferential direction to the mechanical effects that they produce. Even with that simplified view of the crack medium, the effect of the cracks on the strain only has been taken into account.

In the case of Piledriver, which is where this was done first, the material had to be weakened greatly relative to any strength you could ascribe to it from laboratory tests before the observed displacements were obtained. I think it is very clear that the model arrived at has not satisfactorily accounted for the observations.

That is not enough by itself for two reasons. One, there is recovery of about two-thirds of the peak displacement observed. The second aspect of that particular event that is really puzzling is Hardhat, supposedly fired in the same medium, which you are supposed to be able to scale. We concluded that the Hardhat and Piledriver displacements are at least a factor of two different--I think to get beyond where we are now, the next effect that will have to be taken into account is dispersion, again treating the medium as if it were isotropically cracked.

Your study of the effects of closing of cracks in one dimension is a first study in that general direction. We really want to account for a randomly cracked media. In a specific calculation or a specific explosion, when the stress wave reaches the interface between two blocks of material, it just simply isn't perfectly transmitted or perfectly reflected. Whether that partial reflection is important or not certainly depends on the frequency content of the wave that arrives there. If it has an important component at a wavelength like the spacing in cracks, then they are going to be very strong dispersion effects. How strong they are in bursts that interest us, I don't know. But in the Piledriver case there are large cracks, I think on the order of a meter apart. And that is not so short as to be negligible.

At any rate, it appears to me that would be the next logical step. It is different from the usual, well, I would say, thermodynamic view, though it is a general thermodynamic one, of the representation of the stress-strain. It really doesn't fit in quite that framework.

So something other than a simple modification of a stress-strain relation has to be done to properly model the effects of cracks on dispersing waves. We really want to model the detailed interaction of blocks with each other, but I don't think that can even be begun. I don't know whether that would be important, and I hope it wouldn't, except perhaps for the material close in where the dimension of the object

which you are interested in having interact with the soil or the medium is not much greater than the spacing of the cracks. Otherwise, you would hope that you could get some homogeneous or at least isotropic modeling of the effect of the cracks that fail and then you really have a three-dimensional problem to solve and it may be a while before we can do that.

MR. GODFREY: I had reached much the same conclusion of the necessity for adding a dispersive mechanism to the calculations, not just from the cracks themselves. I have been doing a little work recently that I have not seen published before. I took an idealized crack that was in full contact, used a coefficient of friction of dry rock, and solved the elastic wave equations. One finds that even elastic waves hitting this idealized crack at certain angles will not be transmitted completely. So even a perfect crack in dry materials will not reflect at certain angles, much less real cracks.

Thus this reflection of energy is an important thing to include in some way. You also have a dispersion caused by the fact that the materials themselves in a region are not homogeneous. For the Mineral Load shot, which was like 15 tons up at Cedar City, buried 100 ft, the area was extremely carefully selected to be homogeneous. I went into the hole, and one could see, since it was core drilled, the jointing planes fairly clearly and something like three of four were observable up near the top, but in the last 50 ft there was nothing. Now there are two lines of data from this shot. (There were actually four, but unfortunately only two lines were useful.) These gave just completely different results, something like a factor of three in peak velocity.

If one goes back to the drilling and examines closely the cores that were recovered, one can begin to determine a location of discontinuities in material and it is conceivable that one can make these data consistent by changing the model of a certain radius on one of the lines.

My point is that for larger explosions you are going to get variations in properties of materials of a factor of two or so, and this in itself will cause a dispersion. It is conceivable that this kind of thing can be included in a code calculation. One might get 100,000 zones and do sort of a Monte Carlo, taking the extremes of core data for that whole area and then in a random way changing the properties of the materials in regions or even zones. In a very real way the stress waves would be reflected from these changes of impedance. It is not outside the state of the art to do that.

A couple of other comments. Hank Cooper referred repeatedly to spherical decoupling, and I just want to say that God didn't say that cavities had to be spherical for decoupling experiments or that that is the optimum decoupling. We would like to do a little work on what a long tunnel might do, because it might conceivably be much easier

to drill, with the moles now that can chew their way through rock at incredible rates, something like 30-ft diameter. Maybe this is the way to go in decoupling.

I think some studies should be made of nonspherical cavities.

One other thing Hank spoke of, and we got a lot of mileage out of the spherical cavity loaded elastically and what was required in the way of zoning and it made it a good test problem. We have been doing a little work on Lamb's problem which is a point source on a surface of a half space which has been solved analytically. This might be a very nice test problem, not so much for ARPA, I guess, but for DASA, who again may have various code people try to compute Lamb's problem and compare the answers with the analytical solutions.

MR. FRASIER: There is also a solution to the sphere in a half space. It is not given in a very convenient form, you have to use an infinite number of branch line intervals to get back to your record. But this might be a more realistic thing to try than to consider Lamb's problem; Lamb's problem has been done numerically.

MR. ROTENBERG: For one of the graphs you showed, you said that experimental results are consistently higher than the results that came from the computer. Can you be more specific?

MR. COOPER: One of the comments that I made was about a plot of particle velocity and displacement for a granite. You get a difference from the field data of about a factor of two or three between the line you might draw through the particle velocity for granite and tuff. That is consistent with an impedance difference in about an order of magnitude-- ρC^2 is like the confined modulus of the material. So that would be an order of magnitude difference instead of the factor of two for the confined models.

All I am saying is that based on laboratory tests you probably would expect to have bigger displacement in the softer rock. I think this was pointed out in one of the last meetings we had here, because it was about that time we recognized this fact. To me it means that the in situ state is a thing that is really controlling displacement.

If I were Harper, I would look at that curve on the right and wonder a little bit. I wonder if I could really feel good knowing about displacement of rock within the scatter of a factor of three. I don't know.

If joints control things, then maybe it is not the parent material property that is really important. The fact that it is jointed anyhow ....

MR. TRULIO: I don't know if they control, but I think everybody would agree that they are very important and it would be a good question to

ask whether cracked samples of material on which you can experiment in a laboratory would give you correct results for the effects of joints that are present on a much larger scale in the field.

MR. GODFREY: I think that this tuff was probably saturated ....

MR. COOPER: I don't know whether it came right out.

MR. GODFREY: The much more dominant characteristic of the rock that does affect its displacement rather drastically is the porosity.

MR. PHINNEY: This may be changing the subject somewhat, but I would like to inquire about other analytical bases for the calculations. I think there are two examples that would come to mind. One is that you are essentially taking the differential equations necessary to define the problem logically and setting up a certain point. You can also formulate these problems as integral equations for which the numerics are a wholly different set of problems. I tried to solve that problem on a machine which is now no longer in circulation and found that the integral equation approach at that time was in fact very slow.

The other one is this solving of differential equations. I noticed that you had the greatest difficulty in getting the discontinuities to propagate, just as one would expect. In such problems as this, you are not interested necessarily in complete high fidelity of the spectrum from zero to infinite frequency.

I would like to inquire of this general field of analytical measurements as they are being turned into codes.

MR. TRULIO: You have actually raised several questions here. First, with respect to integral formulation, I am not sure in exactly what sense you mean it, and there are at least two. The differential forms of the general conservation laws are not the basic ones. They go actually back to the first or second year of physics. You derive them from more basic forms and there is an important difference between them. The integral statement might be perhaps written as instantaneous rate equations for regions, in which case they don't contain any special derivatives at all and that is the important difference between the two. There is a direct measure of fidelity of representation of those equations. It leads you to approximate solutions, but you can still ask how well the transformations that you can make about them are reproduced by the different equations.

Another level of integral equation is the linear field problems. By using the linear field problem any field can be used to find the strength of sources distributed over surfaces, which takes one independent variable out of the problem. That may be important, actually, for late-time calculations, low-stress calculations, wave propagations in earth media, as you may wind up doing a complete integration

117

of the equation for motion out to a point where the medium or where the solution is like a linear elastic one.

MR. GODFREY: I would like to come back just a second to Lamb's problem because apparently what I said was misunderstood. I want a problem that has been solved numerically, and I am aware that numerical solutions have been obtained. I called them analytical because they came from evaluating a whole bunch of integrals. They are analytical as compared to numerical differences. There are some on surface explosions where there are late-term surface motions that are much larger than any of us have been able to predict.

I think it is important to prove the ability of these codes to compute, for example, Rayleigh waves and surface phenomena. To my knowledge this has not been demonstrated yet. So I say again that Lamb's problem, where a solution is available, would be a very good test problem. The only person I have heard of who says he has done Rayleigh waves successfully is Mark Wilkins, but I haven't yet seen his work.

MR. PHINNEY: This problem has been solved and it is used by seismologists. In fact, the Lamb's problem with several interfaces has been numerically solved. The problem in applying it to this kind of a problem is that the solutions in the dimensions of the source are always small compared to any wavelength involved in the problem. But you don't have to consider how you are going to represent a finite source.

MR. TRULIO: Why not use the source that is produced from the point after a certain time?

MR. PHINNEY: That is one of several things you might do.

The other point is that several groups, in particular Peters and his colleagues, have been taking earth models which are characterized by 500-odd layers to represent the radial variations, putting in earthquake models, and computing exact "responses." I wouldn't say it has come to being a progressive way, but I think that people are interested in earth structure to the point of calculating some useful things.

MR. CHERRY: Hank Cooper, you didn't mention model studies. How do you see model studies as a possible means? I feel that they are an important diagnostic tool in which you take a block of material that you either fabricate or bind and do whatever test you have to do on it. Do some high-explosive testing in that block with appropriate pressure gages and velocity gages and run the corresponding problem on code to see how well you are doing on the stress wave propagation aspect and the equation of state.

MR. COOPER: I am sorry; it slipped my mind. In addition to that, I think the areas of phenomenology that I showed on the last chart should include not only that kind of thing but also tasks where you try to

look at blocky material. There are some basic phenomena that could be understood from the kind of experiments that can be done in the laboratory. I jokingly referred to these sugar tests that we did, and I have some pictures. I think that was a pretty interesting little thing. I am not sure what to do with the data. We are just trying to calculate it. To calculate it we have to use an isotropic material approximation, because that is surely the way it behaves.

MR. HEALY: I would like to generate an argument about something that I tried to start this morning but was overlooked in the heat of the other discussions. Let's look at the explosion at the time when we usually stop looking at it, after it has gone off and a seismic wave has presumably passed away. There we sit with some kind of a cavity. We may have a pressure in the order of perhaps a kilobar at a place where the overburden pressure may be 100 bars.

Has anyone solved this static problem as to whether that hole can contain that pressure?

MR. CHERRY: I don't think it can happen. I think the cavity stops growing when you get to the overburden pressure. Depending on the strengths of the material, pressures like overburden may be a factor. The French case was different. I think in that case we came to--I have forgotten now--maybe a factor of five. But the final cavity pressure was a factor of five above overburden when the cavity stopped.

MR. HEALY: Take Salmon, for instance. Consider the Salmon explosion, which is the one I have the numbers for in my head. The static cavity was 17 m after it had been drilled. Some 4 percent of the energy left in that cavity is either permanent displacement or seismic energy. So you have 96 percent of the energy that really has to be confined in that 17-m cavity. I don't think that is possible at 80 bars displacement. I would calculate the number much larger for cavity pressure, but I don't know anything about these things.

MR. CHERRY: 96 percent is high.

MR. HEALY: You just integrate the displacement of the cavity wall that can account for 4 percent of the displacement of the crack. The other 96 has to still be in that cavity until it has time to pass out through thermal conductivity or some other mechanism.

MR. TRULIO: Or pass through.

MR. STEPHENS: At the time of re-entry into Salmon most could be counted entirely as thermal energy in the profiles around the cavity. I don't remember the exact number, something like 90-95 percent.

MR. HEALY: This took months to come out to that few meters from the cavity. Immediately after the shot that energy was tied up in the gas pressure inside this cavity. How do you handle that problem from there

on? Have you actually calculated and showed that this pressure goes out in cracks? Does it go out in one large crack or thousands of little cracks? What happens at that point?

MR. TRULIO: We haven't made the appropriate calculations. So, at the times you are talking about, no.

MR. GODFREY: I don't believe your 96 percent, because to get that number you have to know the pressure-time history of the cavity. I don't know how you get that from your model. How did you determine the 4 percent?

MR. HEALY: It is the energy that actually comes out into the elastic zone. I will give you 10 percent.

MR. GODFREY: There is a lot more plastic energy dumped into the region out to the elastic radius, a tremendous amount.

MR. HEALY: I don't know anything about these; I am asking you that. This is the whole containment question. Now if you can't tell what happened to these fractures, we can never begin to know that a nuclear explosion is contained. I put it to you as an important problem. You look at these pressures, and from a simple calculation you arrive at many times the expected overburden pressure. I have never seen any calculation or any evidence that demonstrates that this is not true.

MR. TRULIO: I don't know of anyone who has carried a calculation out to a minute or hours from a kiloton, let's say. But there has been some thought given to it.

How does the heat that might otherwise be in the gas and contribute to a large pressure or a much greater one than the overburden get lost from the gas? Well, if the cavity explosion is a nuclear one the gas will be very hot. You can imagine simply boiling off and popcorning, I think, is the word now, pieces of material from a wall which will then start traveling through the cavity, small pieces of material that will go a lot faster than an ideal model of a smooth wall with whatever the low conductivity of salt is. The cooling might be a lot faster, especially since these materials are wet; we can cause the water to boil and break up into small pieces at the wall. In addition to that, the roof of the cavity usually collapses and there is some quenching from that.

But I don't know of anybody having calculated the things I am talking about in detail.

MR. HEALY: You see, there is another plausible explanation as to why this may persist, and that is if you have a single large hydrofracture and this material can move out. It is conceivable that you can have vertical fractures--single, large vertical fractures--which open up and take the gas. This would have a very profound implication for a number of problems that are important.

120

MR. GODFREY: Someone else had tried to say that. It is my understanding that for a wide range of underground nuclear explosions, the models of the explosion with the pressure carried down to that of overburden pressure together with the code run that predicted the size of the cavity compared with the measured size of the cavity, gave extremely good correlation. People have not found large vertical cracks.

MR. HEALY: Wait a minute. Now let's not get mixed up here. The cavity should be directly predicted by the three-dimensional code, because it is controlled by phenomena that are propagating faster than the usual rupture velocity for fractures. So the formation of the cavity is perfectly plausible.

The second point is that examination of the test site reveals a lot of large vertical fractures that aren't scattered all over the surface and these are not specifically symmetrical, so the evidence to the contrary would suggest that indeed vertical fractures are formed by processes which we have not yet explained.

MR. GODFREY: But these fractures don't, as far as I know, go near the cavity; they are just surface.

MR. HEALY: I don't know of any experiment where they have drilled to look for these fractures. They are hard to find with a single vertical drill hole.

MR. GRINE: If they went through the cavity, they wouldn't be hard to find; they would leak radiation everywhere, and they don't.

MR. HEALY: You mean there is no radiation leakage?

MR. GRINE: Some of them have some, but most don't.

MR. HEALY: It is true that most haven't leaked to the surface. I am not saying that I understand this phenomenon. I am saying that I don't understand the physical observations, and I haven't found anyone who has done a calculation that helps me to understand it.

MR. GRINE: The calculations these gentlemen are talking about do include heat depositions behind the stress wave that propagates out from the fracture. That is what Godfrey was talking about when he said there is a tremendous amount of energy deposited by the plastic wave and most of the calculations I have seen give very large cracks. In what sort of percentage area do you get deposits?

MR. TRULIO: Within one cavity diameter?

MR. GRINE: Yes.

MR. TRULIO: If it is compactible you may get half.

121

MR. GRINE: Tuff?

MR. TRULIO: Tuff, yes.

MR. TRULIO: In the region of plastics identified in these calculations, the time scale that hasn't been calculated is the time that would be required for thermal conductivity to be important. Now it is possible that calculations can be made which don't include motion; the inelastic cooling is assumed to represent radial cracking and the extent of that could be translated into a void space available to hot gases in the cavity, then a conduction calculation could be made. But I am not sure that procedure would give a valid answer. These cracks may not really be open to the gas.

MR. GODFREY: I noticed in Hank Cooper's presentation that for the rock model we think most appropriate, almost no tangential stress ever built up. In fact, because the radial compression tends to extrude the material, if you want to look at it that way, you essentially get very little tendency for any of these radial cracks to form.

As I say, all the calculations on the underground things don't even require cracking to be able to predict the cavity size.

MR. HEALY: But this, again, isn't really relevant to what I am saying. If there is no tectonic stress and the problem is truly one dimensional then what you are saying is correct. But if there is tectonic stress, a single crack correction would dramatically change the physics of the problem. The mere fact that you do, in cases where we have seen the cavities, get approximately the right size, I don't think is convincing. There are others that are suspiciously deformed, particularly if you look at the whole set of strain measurements around an explosion.

MR. GODFREY: My only comment would be that tectonic stresses aren't likely to be very much greater than the overburden stresses, and they are sort of trivial in the vicinity of the cavity to the forced stresses to which they are subjected.

CHAIRMAN SIMMONS: I wonder if there are comments from the audience.

MR. RINEY: I would like to come back to the discussion relative to the fact that we do not have a competent rock surrounding the cavity and also some remarks by Chuck and his cohort that maybe we should model material in such a way that there is some attenuation involved in the modeling itself. Relative to this it might be good to recall some of the rock mechanics that we had this morning in which the importance of both the presence of pores and the presence of water in the pores was found to have a very strong effect on the mechanical behavior, and it is also known to have a very strong effect on the wave propagation in the pores.

For the last year we have been working on a program which does address some of these questions and I think some of it might be relevant to the discussion. So far as the calculation similar to what Hank Cooper mentioned, instead of having cracked media we were considering alternating layers of different materials. It turned out we were considering tuff and water in our layers. It is very interesting that in compressible materials we found the attenuation was greatly decreased by the presence of the interfaces as far as a short pulse. Apparently the rarefaction wave which attenuates the pulse is not as effective in this case. This, I think, is what you were also saying.

Another aspect that probably is of more interest is that we thought it might be useful to look at the interacting continuum whereby we can consider the presence of pores and water and the matrix material all within a continuum context. This has some advantages insofar as one can get dispersion relationships in this theory, which is, I think, the sort of thing you were addressing your remarks to earlier. This type of modeling we are currently doing for tuff both completely and incompletely saturated. We are also in the process of making some wave calculations using a numerical scheme.

MR. HARPER: I would just like to make an observation as a seismologist.

The yield versus magnitude curves are indistinguishable for unsaturated tuffs, granite, and salt within the scatter of the data. They all lie on the same curve. In addition, when you go to a yield of 100 tons or less, even the alluvium values are given the same amplitude. At times I suspect that the detailed labor on rock mechanics is somehow missing a point here, because all of the materials are looking alike. The details of the fractures and the details of this and that seem to be somehow missing the point, because they are all coming out the same way. I don't know why that is so. I don't play any of your games. Since it is all about the signals anyway, I think you ought to keep remembering these seismological facts of life.

MR. CHERRY: Is there anything unique about the 3 cps part of the spectrum that you are looking at? I don't know if anybody has looked at coupling in terms of frequency content or not.

MR. HARPER: The uniqueness is that we have to look there and to minimize the effect of natural background noise ....

MR. CHERRY: This is first arrival.

MR. HARPER: That is correct, these are first arrivals.

MR. CHERRY: You haven't looked at anything else on the record to see if maybe you could distinguish between the three media, say in terms of the Rayleigh wave or in terms of other frequencies?

MR. HARPER: No. Well, there is another story about the Rayleigh waves. But for the shots at NTS the two will go up and down together. If the P signal goes up the Rayleigh wave will go up. And if the P signal goes down the Rayleigh wave goes down. So on a plot of Rayleigh wave versus P wave, ignoring the medium, they all lie on a very narrow curve and you can't tell them apart on a ratio of Rayleigh waves to P waves. From your point of view, why we do it is highly arbitrary, and in a way it is, but again it is living in the facts of seismology. There is not much energy of higher frequency content, and if you are going to go a bit longer period than one second, then you start getting in trouble with large microseismic influences that are in the ground.

MR. TRULIO: You say they all look the same. You mean the same yield tamped in all these media give almost the same teleseismic signal?

MR. HARPER: That is right.

MR. TRULIO: Within what factor?

MR. HARPER: I think it is less than two.

MR. TRULIO: I don't think the conclusion of that fact is that we go home; I think there are still some more questions to ask.

MR. DRAKE: For the last six or eight months, I have been trying to model joints or cracks in a rock by statistical methods, and I distribute these things, assuming that they are random, isotropic, homogeneous, and so on. These joints have a certain delay property. I get some very interesting results from this very simple model in that I find a rise time that becomes proportional to a travel time up to a certain distance and find that this distance increases as if it is a square root distance. It seems possible to me that at a certain range, the statistical property of this material can take over and make it all indistinguishable from one shot to the other if you can go far enough. The statistical property tends to work on this pulse unless they all turn out to be the same.

MR. MADDEN: What is the time scale of the calculations that are being referred to here? I was once involved in some electrical measurements connected with the pressures of setting up fluid flows and they were talking about a time scale of days. The seismologists are concerned about time scales of seconds.

MR. TRULIO: A second is the time scale of the experimental calculations, mainly because we have communicated enough to realize that roughly, the 1 cps kind of frequency is the one you are interested in. So the calculations are carried that long.

CHAIRMAN SIMMONS: I expect that our 5 o'clock cutoff time is inviolable for this group of individuals. I would like to keep the last half hour for the discussion of ILLIAC. Maybe we should turn now to David McIntyre and talk about the ILLIAC for a while.

THE ILLIAC IV COMPUTER

David E. McIntyre
University of Illinois

This introduction for ILLIAC IV is written for a computer user who has only an acquaintance with the hardware involved in a digital computer. For a more complete description consult references (1) and (2)*.

A stereotype computer can be characterized using the boxes shown in Figure 49a.

STEREOTYPE

CONCURRENT
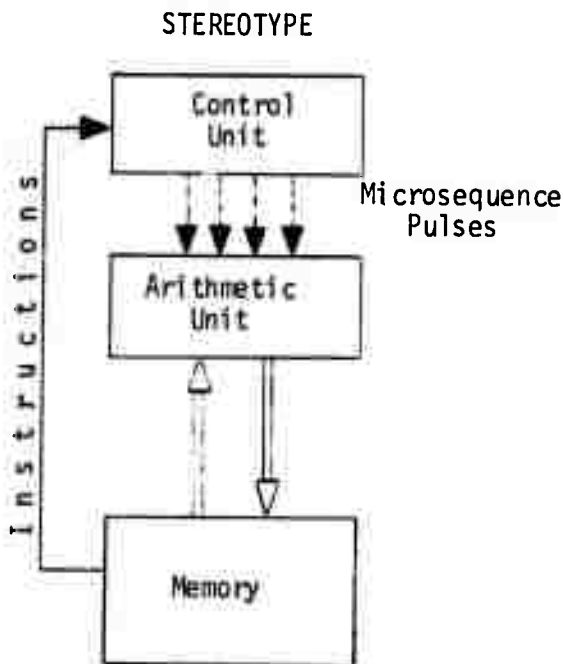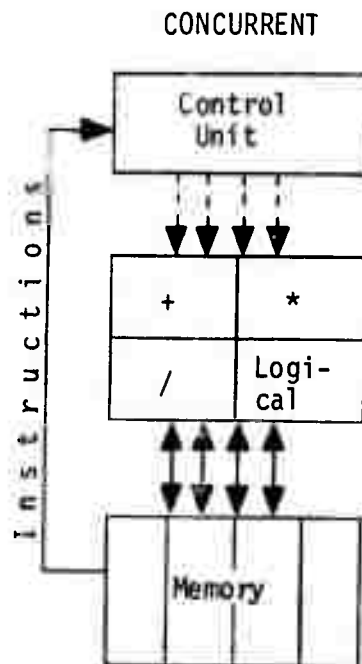
Figure 49a.

Figure 49b.

It is composed of a memory which holds operands and instructions; a control unit which fetches instructions from the memory, decodes them, and issues control signals (microsequence pulses) that operate, or drive, the arithmetic unit. The arithmetic unit performs the computational operations: addition, logical operations, and multiplication, on operands that have been supplied from memory, and returns the result

---

* References are listed numerically on page 143.

to the memory. In effect, the control unit monitors and controls the flow of information between the memory and the arithmetic unit in addition to actually operating the arithmetic unit.

A typical sequence of events that takes place during operation is:

1. An instruction is fetched from memory to the control unit;

2. When it arrives in the control unit, it is decoded;

3. If the instruction involves operands in memory, the operands are fetched to the arithmetic unit;

4. When the operands arrive in the arithmetic unit, the computation (for instance, subtraction) is initiated and monitored by the control unit until complete;

5. After completion, the result is stored to memory.

There are several ways one can modify this stereotype design to achieve an increase in computing speed. One way would be to add an additional control unit and arithmetic unit. This is the multi-processor configuration. Another way (Figure 49b) would be to divide the arithmetic unit into a group of functionally independent subunits, each of which could be operated independently by the control unit. As the control unit decodes instructions, it will determine when two or more consecutive instructions use separate functional units and are independent of each other. If this is the case, the separate instructions are allowed to proceed concurrently in the separate functional units rather than sequentially, as was the case with the stereotype machine. This is the approach employed in the CDC 6600 and IBM 360/90 series. Of course, if the number of operands which can be processed by the arithmetic unit is increased, the speed with which the arithmetic unit can obtain and store operands from the memory must also be increased to avoid a bottleneck. This can be achieved by partitioning the memory into several sets of memory banks. A memory operation can then take place simultaneously in separate memory banks.

Figure 50 shows how the stereotype design has been modified in the design of ILLIAC IV. This figure describes one quadrant, or one-fourth, of the ILLIAC IV array. The control unit operates in very much the same manner as the control unit in the stereotype computer. Instructions are fetched from the memories to the control unit where they are decoded and microsequence signals are produced. The microsequence signals are duplicated 64 times, and each set of microsequence signals is passed to a separate arithmetic unit. The same set of signals operates 64 different arithmetic units and increases the number of arithmetic operations that can be performed by a factor of 64. An arithmetic unit is referred to as a "processing element" (PE). Each arithmetic unit (PE) can fetch or store operands only to or from its own unique memory bank. The control unit, however, can fetch instructions from any of the 64 memory banks. The restriction that each arithmetic unit performs memory

operations only with its unique memory solves some problems and poses some others.
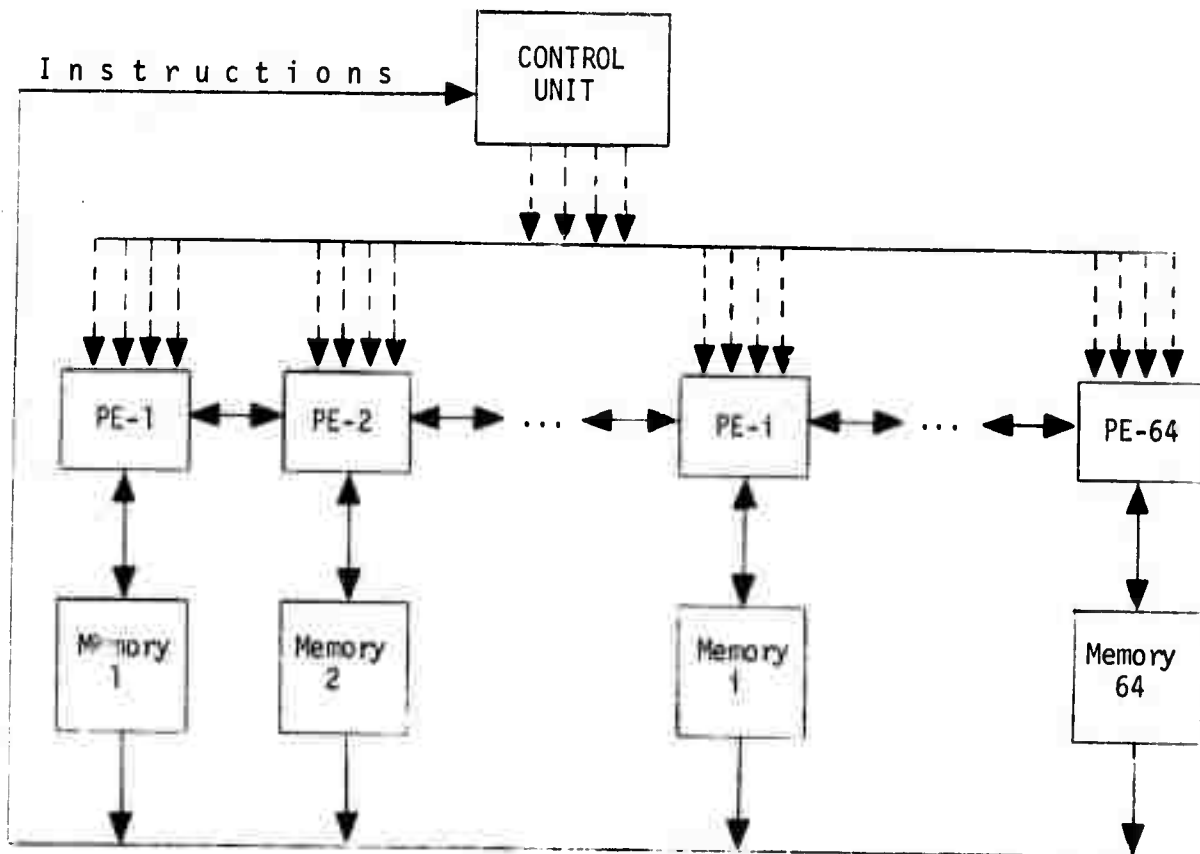
Figure 50.

If the memory banks of ILLIAC IV were arranged with a large crossbar switch so that any PE could access data from any memory bank, there would be delays imposed because of the distance that the signals would have to travel. Furthermore, if PE-i required a datum that was stored in memory k and, at the same time, PE-j also required an operand in memory k, PE-j would be forced to wait a complete memory cycle time before it could receive its operand, since only a single memory operation can take place using the same memory bank at one time. These delays are referred to as bank conflicts and are encountered even on conventional designs like a concurrent computer. Therefore, by assigning each processing element its unique memory near the PE, signal line delays can be minimized and the possibility of bank conflict will be eliminated since only one PE will be making demands on a given memory bank.

Programs run most efficiently if all the operands used by PE-i can be stored in memory i. This is a very restrictive condition and is not always possible. Occasionally PE-i needs to use an operand that is stored in PE-j's memory. This is accomplished by fetching the operand from PE-j's memory to the operating registers in the arithmetic unit and then transferring it to PE-i. This process is called routing and will be explained in more detail later.

Processing Element

The processing element in the ILLIAC IV is basically a four-register arithmetic unit. (See Figure 51.)

Figure 51.

128
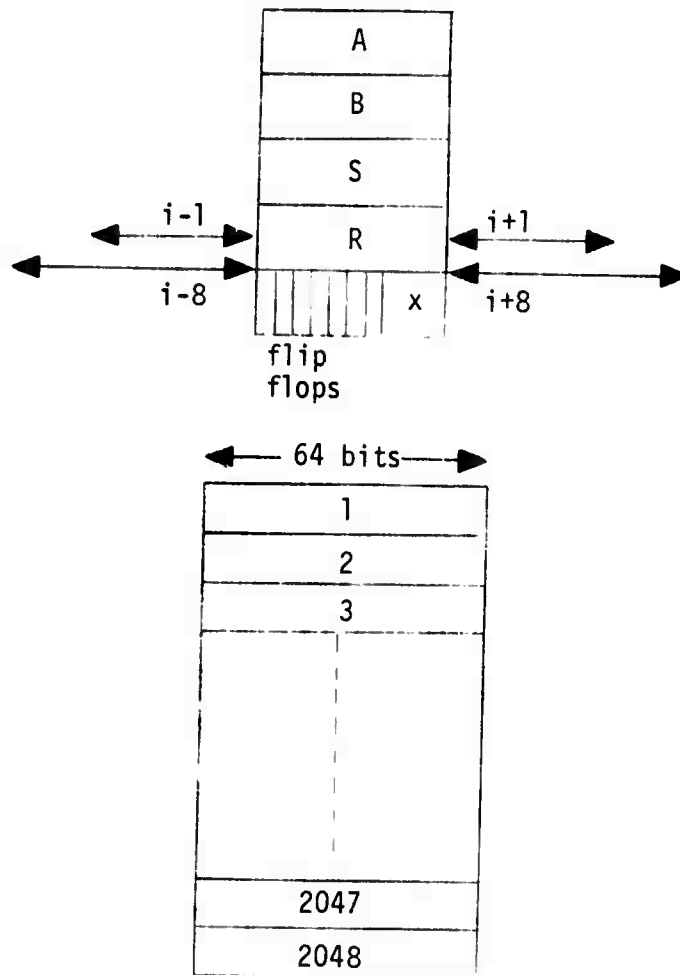
There is an A register and a B register, used to hold the operands for arithmetic and logical operations. Operands for arithmetic operations are placed one in the A register and one in the B register. The operation is performed and the result left in the A register. The S register is provided as a kind of scratch-pad memory to avoid making repeated accesses to memory to fetch or store intermediate results. The R register is used to transfer information among the PE's in the routing operation. Each of these registers is 64 bits long.

The R register in PE-i is wired directly to the R register in PE-i±1 and PE-i±8. The routing operation uses the R registers and can be visualized by considering the 64 R registers as a large 4096 bit register. Upon executing the command "route 1 to the right," this long register is shifted end around 64 bits to the right. (See Figure 52.)

64 bits
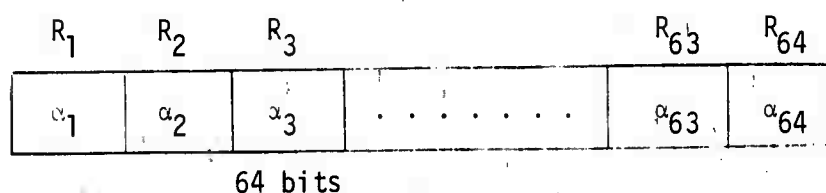
Route 1 Right Accomplishes
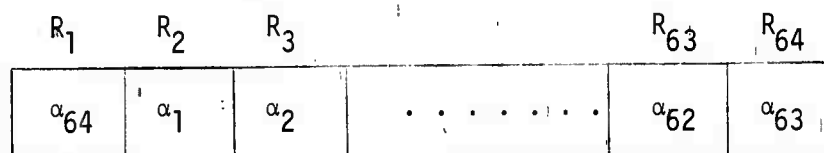
Figure 52.

Routes can be performed toward the right (in the direction of increasing PE number) or left. A distance-8 route (shift of 512 bits) is provided so that information can be rapidly sent between PE's with greatly different numbers. Displacements of ±1 and ±8 require about 100 nsec. Arbitrary distance routes are decomposed by the hardware into several consecutive routes of distances 8 or 1.

There is also an 18-bit index register (x register), which is used mainly to increment a basic memory address. Finally, there are eight 1-bit flip-flops which can be used to store the results (true or false) of tests, logical operations, etc. Each PE memory is composed of 2048 64-bit words. It is a semiconductor memory with a cycle time of roughly 200 nsec.

It should be pointed out that since the microsequence stream controlling PE-i is exactly the same stream controlling PE-j, the PE's

are constrained to execute exactly the same instructions at exactly the
same time. When PE-1 is performing an addition, PE-5 cannot perform a
multiplication. There are two degrees of local autonomy provided for a
PE. The first degree of autonomy involves "turning off" or disabling a
PE. A disabled PE can perform no operations. A PE can be disabled
either on command from the control unit or as a result of some con-
ditional test. For instance, at the end of an arithmetic operation, the
control unit can issue a command which is interpreted as "any PE that
has computed a negative result, turn yourself off." Once a PE is turned
off, it can no longer turn itself back on and must be enabled on command
from the control unit. The other degree of independence available is
that each PE may use a different memory location for a memory operation.
This is accomplished by incrementing a base address by the contents of
the index register in each PE. Suppose PE-17 is to store the contents
of its A register in memory location 35, while PE-18 is to store the con-
tents of its A register in memory location 45. The index register in PE-
17 would be set to zero, and the memory index register in PE-18 would be
set to 10 and the control unit would issue the instruction "store to
location 35 incremented by the index register." In PE-17 the memory
would be incremented by zero, and the store would occur to location 35.
In PE-18, the address 35 would be incremented by 10, and the store would
be performed into location 45. These two degrees of freedom associated
with each PE actually provide a great deal of flexibility in programming
ILLIAC IV.

PE's can be operated either in 64-bit mode or in 32-bit mode.
In the 32-bit mode, each 64-bit is considered as two 32-bit words, and
two 32-bit floating point operations can be performed in roughly the
time required to perform one 64-bit operation. In 64-bit mode, floating
point numbers have 48-bit mantissas, leaving 16 bits for exponent and
sign. In 32-bit mode, mantissas are only 24 bits long.

Operation Speed

Table 1 compares the execution times for common operations
between a single processing element and another high speed digital
computer, in this case the CDC 6600. In all fairness, it should be
pointed out that there is often a great deal of concurrency obtained
using the 6600, that is, several separate floating point operations can
be going on at one time. There is also a limited amount of concurrency
in ILLIAC IV. If two consecutive instructions to be sent to the PE's
are independent and do not require the same components in the PE, they
may be executed concurrently. Acknowledging that this is rather a rough
comparison, it is fairly reasonable to equate in floating point com-
puting power a single PE and one or two 6600's. If all 64 PE's are
enabled and doing useful work, they can produce floating point operations
at a rate comparable to between 64 and 128 CDC 6600's.

Table 1. Comparison of Execution Times

	PE	6600
Memory to Operating Register (fetch)	350 ns[a]	800 ns
Floating Add	250 ns	600 ns
Floating Multiply	450 ns	1000 ns
Floating Divide	2750 ns	3000 ns
Register-Register Transfer	50 ns	300 ns
Operating Register to Memory (store)	300 ns	1000 ns

[a] $ns = 10^{-9}$ sec

Control Unit

Figure 53 gives a functional representation of the major components in the control unit (CU).

There is a local data memory composed of 64 words. This local memory can be filled from any location in any PE memory and also stored to any location in any PE's memory. There is a block of 64 words called the Program Look Ahead (PLA). This block of words provides an instruction queue, and its operation will be explained in a later paragraph. The arithmetic unit in the control unit is a very simple unit and is restricted to performing logical operations and fixed-point addition and subtraction, obtaining operands and storing results only within the local data memory.

The instruction decoding logic decodes instructions provided from the program look ahead. If the instruction is an instruction to be executed by the array of processing elements, the decoded instruction is fed into the microsequence generator where the microsequence pulses are generated and sent down control lines to drive the processing elements. If the instruction is one to be executed in the control unit, the decoded instruction is issued to the simple arithmetic unit. Most of the instructions executed by the control unit involve housekeeping operations associated with loops or indices. These housekeeping

131

instructions can be executed concurrently with arithmetic instructions
fed to the PE's. (This concurrency is not related to the concurrency
in PE instructions, which was mentioned previously.)

CONTROL UNIT
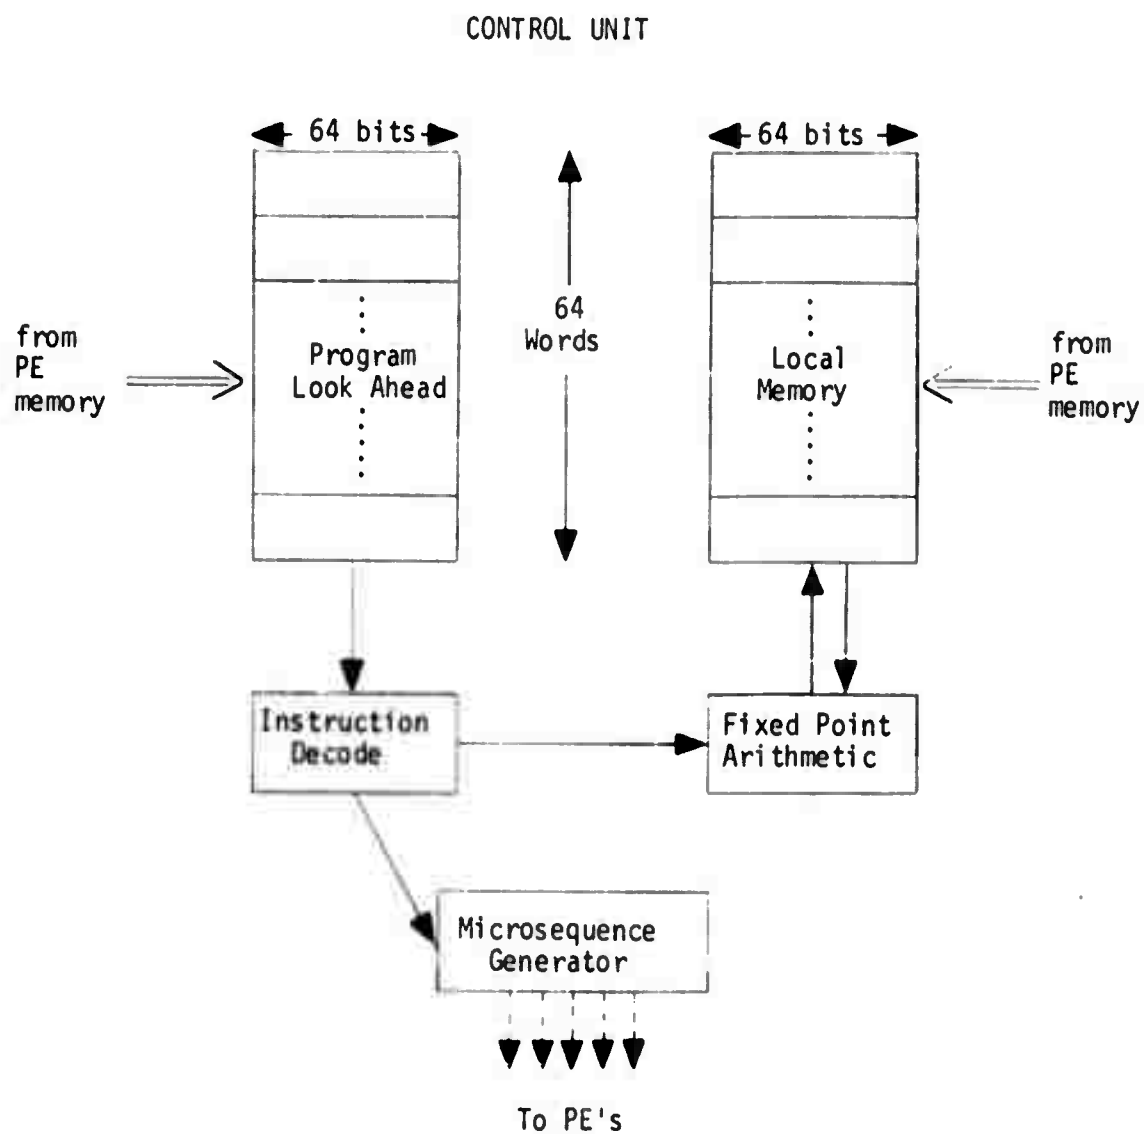
Figure 53.

Operation of Program Queue

An ILLIAC IV machine language instruction is composed of 32 bits. The 64 words (each word is 64 bits) in the program look ahead provide a queue of 128 instructions. Loops containing up to 128 instructions can be executed without any reference to PE memory. The 64 words are divided into eight sections of eight words each. When the control unit is executing the instruction contained in the fifth word of the eight-word section of instructions, it checks to see if the next eight words of instructions are already contained in the PLA. If the next eight words are not contained in the PLA, it will issue commands to bring the next eight words of program to the PLA and destroy the oldest subsection of eight words. This effectively eliminates many of the delays imposed by instruction fetching, except for the case when a jump is made to a section of the program that is not contained in the PLA. For a large range of programs that have been simulated, it has been found that the control unit is delayed waiting instructions to be fetched from memory much less than 1 percent of the time.

Control Unit Processing Element Communication

Operands and control information can be transferred between the control unit and the PE's in several ways:

1. The control unit can broadcast a 64-bit word to all PE's simultaneously. The word originates in the local data buffer, or the arithmetic unit in the CU, and the destination can be any of the 64-bit operating registers in the PE.

2. The control unit can broadcast a 64-bit word with one bit going to each PE, that is, bit one would go to PE-1, bit two to PE-2, ... bit 64 to PE-64. The destination of the bit going into a PE can be any of eight 1-bit registers in each PE. This is a method by which PE's are enabled and disabled. For instance, if it is desired to enable all even numbered PE's and to disable all odd numbered PE's, the control unit constructs (using its arithmetic unit and logical operations) a 64-bit word which has alternating ones and zeros. This word is then sent to the microsequence generator where one bit is sent to each PE, disabling all PE's that receive a zero, and enabling all PE's that receive a one. The 1-bit register that specified whether a PE is enabled or disabled is called the mode register.

3. The control unit receives information from the processing element in the reverse of the method previously described, that is, a bit is sampled from the 1-bit mode register in each processing element and assembled into a 64-bit word in the control unit. The control unit can use this facility to determine which PE's are enabled by assembling a 64-bit word from the single bit mode registers in 64 different PE's.

133

4. The control unit can fetch words from any PE's memory into the local data memory or into the program look ahead. The fetch can consist of a transfer of one 64-bit word or a transfer of eight contiguous 64-bit words. The fetch of eight contiguous words requires only slightly longer than the fetch of one word, thus is a high speed method of getting large amounts of data into the control unit from the PE memory. All of the fetching to the PLA is automatic, as was previously pointed out.

The ILLIAC IV System

Figure 54 shows the ILLIAC IV system organization.

TAPES, READERS, CRT'S, PUNCH
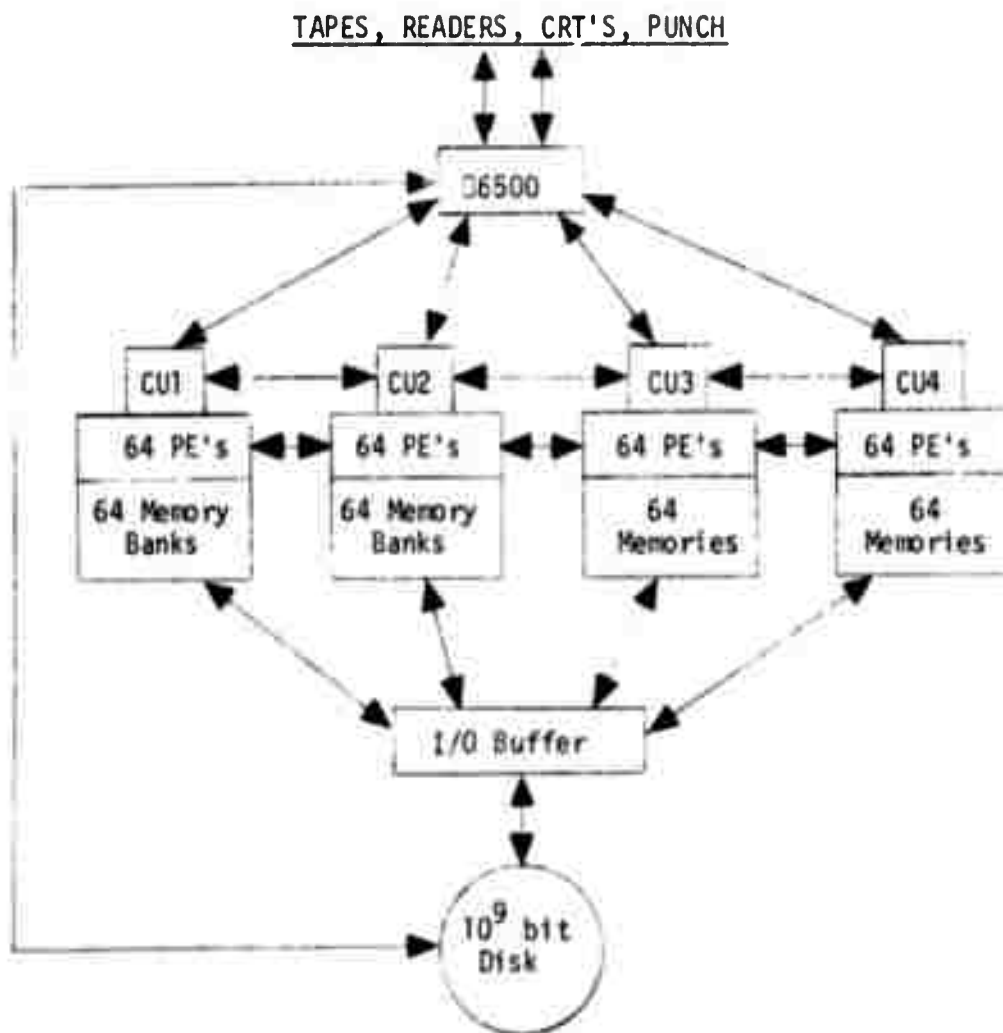
Figure 54.

It is composed of four identical control units (CU's), each control unit driving 64 PE's with 64 PE memories. The CU's are con-nected by lines which allow all to execute exactly the same instruction stream. In this "united" mode of operation, routing is provided across quadrants and end around from PE-256 to PE-1.

ILLIAC IV can be operated in several configurations. For example, all control units can be executing the same instruction stream, or each could be executing a different instruction stream; also, two control units could be executing one instruction stream and two exe-cuting another. It is possible to change the configuration during the execution of a program, but it is felt that this is not an extremely practical facility and does require certain careful programming con-siderations.

ILLIAC IV communicates with the outside world through the Burroughs B6500 computer. The B6500 is very similar to the B5500 but is essentially five times faster. The data base for programs which are not core contained resides on a 10^9 bit head-per-track disk. This disk has two controllers, and each controller is capable of transferring into or out of the ILLIAC IV memory at the rate of 500×10^6 bits per second. If input and output were being carried on simultaneously using both con-trollers, the effective transfer rate can be 10^9 bits per second. This disk has a revolution time of 40 msec, giving an average access time of 20 msec.

The average effective access time can be decreased consider-ably below 20 msec when several I/O requests can be accumulated in the I/O controller. There is a mechanism in the I/O controller which compares the beginning disk address of all I/O operations in a queue of requests with the address of the section of the disk that is passing under the read-write heads. As soon as a match is found, an I/O opera-tion is initiated. For example, suppose two I/O descriptors reside in the descriptor queue with the lowest (oldest) descriptor, descriptor a, referencing a disk address that is located 270 degrees away from the disk address that is passing under the read-write heads; and the second I/O descriptor, descriptor b, requiring a disk address that is located only 90 degrees away from the disk address under the read-write heads. (See Figure 55.) The logic in the I/O controller would initiate the I/O operation b first, requiring only a disk rotation of 90 degrees, or a latency of only 10 msec, instead of initiating the I/O operation a, which would require a disk rotation of 270 degrees, or 30 msec latency time. The queue in the I/O controller can contain 24 I/O descriptors.

The B6500 actually exercises control over the CU's and all of the interactions between the disk and the computing array. The control units request I/O of the B6500, and it coordinates all I/O requests and initiates all I/O transfers between the disk and the array. At the end of an I/O transfer, it signals the control unit that the transfer is complete. In addition to performing this supervisory capacity it also does all of the compiling for programs to be executed on ILLIAC IV.

135

All external data used by the ILLIAC IV array goes first to the ILLIAC IV disk. For instance, if an executing ILLIAC IV program needs to read a tape unit, it makes the request of the B6500; the B6500 reads one of its tape units and writes the result onto the ILLIAC IV disk and then initiates the operation to bring the record from the ILLIAC IV disk into the ILLIAC IV memory.

Read-Write Heads
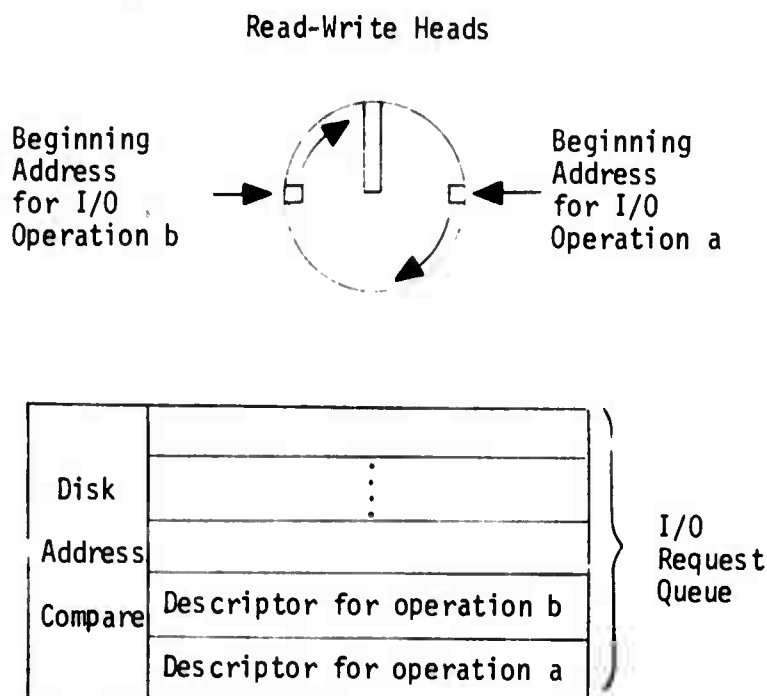
Figure 55.

Programming ILLIAC IV

It was mentioned previously that the separate memories assigned to each processing element provide operands at a very high rate to each processing element, but it does cause a complication when programming ILLIAC IV. The programmer must arrange a storage allocation scheme so that when PE-i needs a datum of information, it is "easily" accessible to PE-i. This does not necessarily mean that the information must be stored in PE-i's memory, since the route instruction makes it possible to obtain operands from memories other than that associated with PE-i. However, it does mean that the transfer of operands from one memory to another using the routing instruction must be done in some regular fashion.

136

Suppose PE-7 needs a datum of information that is stored in PE-2's memory at the same time that PE-12 requires a datum of information that is stored in PE-7's memory. This requires a parallel operation because both PE-2 and PE-7 can fetch operands from their memory and simultaneously route a distance 5 to the right (by performing a route 8 to the right and 3 consecutive routes of 1 to the left), making operands available to PE-7 and PE-12. However, if PE-7 needed the information from PE-2, and PE-12 needed the information from PE-19, this transfer of information would be impossible to implement in parallel because one involves a route to the right while the other involves a route to the left.

Programming ILLIAC IV involves a rather alien situation for the programmer who is accustomed to the conventional machine. He must not only think of some way to implement his mathematical algorithm, but he must also think of a memory allocation scheme for storing his data which allows him to implement the algorithm in a parallel fashion. Memory allocations are not totally new to the FORTRAN programmer. DIMENSION, COMMON, and EQUIVALENT statements in FORTRAN are, in fact, a kind of degenerate memory allocation program.

It is often a common programming trick to reference a doubly dimensioned variable as a singly dimensioned variable. For instance, if A was dimensioned 10 x 10, a smart programmer will sometimes speak of A(27), knowing that he actually is referring to A(7, 3). Occasionally it is expedient to write a subprogram with arguments that are singly dimensioned arrays but to call, or enter, that subprogram using an argument which is doubly dimensioned.

To use the EQUIVALENCE and COMMON statements effectively and employ such simple programming tricks, the programmer must know how arrays are stored or distributed in memory. When using ILLIAC IV, the programmer must consider the memory allocation while he is composing the program rather than simply regarding it as the casual "after-the-fact" problem that it is when programming a conventional machine.

Perhaps it is worthwhile to illustrate the problem of memory allocation and also to demonstrate how parallelism can be achieved by narrating the steps that are programmed to perform a matrix multiply. For this example ILLIAC IV will be constrained to use three PE's. The first step is to arrange to store the elements of the 3 x 3 matrices, X, Y, and the product matrix Z (we will compute X * Y = Z) in the memories as is shown in Figure 56. This form of storage is commonly used for doubly dimensioned variables and is referred to as "straight storage." X_{12} is stored in location 10 in PE-2's memory, Y_{33} is stored in location 27 in PE-3's memory, etc.

	PE-1	PE-2	PE-3
loc 1			
	o o o	o o o	o o o
loc 10	X_{11}	X_{12}	X_{13}
loc 11	X_{21}	X_{22}	X_{23}
loc 12	X_{31}	X_{32}	X_{33}
	o o o	o o o	o o o
loc 25	Y_{11}	Y_{12}	Y_{13}
loc 26	Y_{21}	Y_{22}	Y_{23}
loc 27	Y_{31}	Y_{32}	Y_{33}
	o o o	o o o	o o o
loc 102	Z_{11}	Z_{12}	Z_{13}
loc 103	Z_{21}	Z_{22}	Z_{23}
loc 104	Z_{31}	Z_{32}	Z_{33}
	o o o	o o o	o o o
loc 2048			

Figure 56.

Step 1: Copy the first two of X to the CU local data buffer (LDB). (The contents of memory location 10 in the first three PE's.)

Step 2: Simultaneously fetch the first row of Y to the B registers. Each PE fetches from location 25.

Step 3: Broadcast X_{11} from the LDB to the A registers of all PE's. The contents of the registers are shown in Figure 58a.

Step 4: Multiply and store contents of A in the S register.

Step 5: Fetch the second row of Y to the B registers. Each PE fetches from location 26.

Step 6: Broadcast X_{12} to all A registers. Contents of A, B, and S registers are shown in Figure 58b.

Step 7: Multiply ($X_{12}*Y_{21}$ is formed in PE-1's A register, while $X_{12}*Y_{23}$ is formed in PE-3's A register.)

Step 8: Add contents of S to contents of A and store results in S.

Step 9: Fetch third row of Y to B registers. Each PE fetches from location 27.

Step 10: Broadcast X_{13} to all A registers. (See Figure 58c.)

Step 11: Multiply.

Step 12: Add contents of A to contents of S. Figure 57 shows contents of A registers after addition. Note that the first row of the product matrix has been formed, Z_{11} in PE-1, Z_{13} in PE-3.

PE-1	PE-2	PE-3
$X_{11}*Y_{11}$	$X_{11}*Y_{12}$	$X_{11}*Y_{13}$
+	+	+
$X_{12}*Y_{21}$	$X_{12}*Y_{22}$	$X_{12}*Y_{23}$
+	+	+
$X_{13}*Y_{31}$	$X_{13}*Y_{32}$	$X_{13}*Y_{33}$

Figure 57.

Step 13: Store contents of A registers to first row of Z. All PE's store to location 102.

139

```
                        CU        LDB
                  ┌──────────┬──────────┐
                  │   X₁₁    │   X₁₂  │  Y₁₃  │
                  └──────────┴──────────┘
```

	PE-1	PE-2	PE-3
A reg	X_{11}	X_{11}	X_{11}
B reg	Y_{11}	Y_{12}	Y_{13}
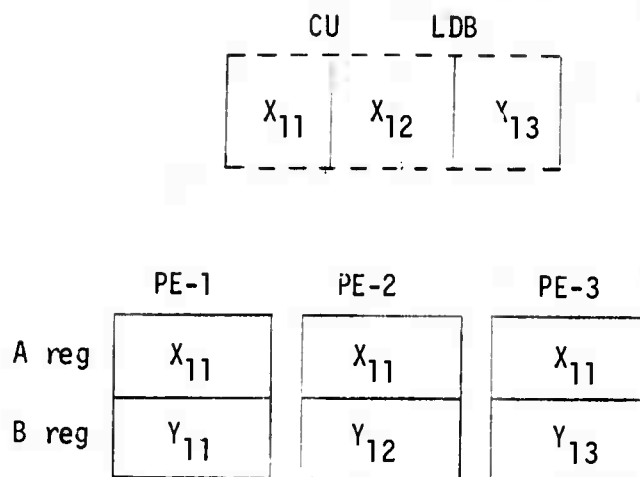
Figure 58a.

	PE-1	PE-2	PE-3
A	X_{12}	X_{12}	X_{12}
B	Y_{21}	Y_{22}	Y_{23}
S	$X_{11}*Y_{11}$	$X_{11}*Y_{12}$	$X_{11}*Y_{13}$

Figure 58b.

	PE-1	PE-2	PE-3
A	X_{13}	X_{13}	X_{13}
B	Y_{31}	Y_{32}	Y_{33}
S	$X_{11}*Y_{11}$ $+$ $X_{12}*Y_{21}$	$X_{11}*Y_{12}$ $+$ $X_{12}*Y_{22}$	$X_{11}*Y_{13}$ $+$ $X_{12}*Y_{23}$

Figure 58c.

140

Now, in order to compute the second row of the product matrix, the second row of X is fetched to the LDB in the CU, and the process is repeated.

In fact, what we have demonstrated through the description of a matrix multiply is that ILLIAC IV can do elementary row operations "in parallel." However, there will be many applications in which it is desirable to perform column operations in parallel, as well as row operations. In particular, matrix inversion and numerical solution of partial differential equations require this facility. An alternate method of memory allocation, called "skewed storage," permits row and column operations to be performed in parallel.

Figure 59b shows the skewed storage technique for a 4 x 4 matrix A. As in straight storage, the first row is stored across the PE's at some location ξ in the PE memory. The second is then "skewed" or rotated once to the right, so that a_{21} is stored in PE-2 (instead of PE-1 as would have been the case with straight storage). The third and fourth rows are skewed two and three PE's to the right, respectively.

Now in order to perform row operations involving the third row, the contents of memory location $\xi + 2$ is simultaneously copied to the PE operating registers. To perform a column operation involving the first column, the index registers are loaded as shown in Figure 59a. The index register in PE-1 is loaded with 0; the index register in PE-3 is loaded with 3. Then, on the command "fetch from loc ξ incremented by the index register," the first column (circled elements in Figure 59b) is simultaneously copied to the PE operating registers. To fetch the second column, the contents of the index registers are simply rotated one PE to the right (dotted portion of Figure 59a), "fetch from location ξ incremented by index register" is executed, and the second column (triangles) is obtained.

	PE-1	PE-2	PE-3	PE-4
X register	0	1	2	3

3	0	1	2
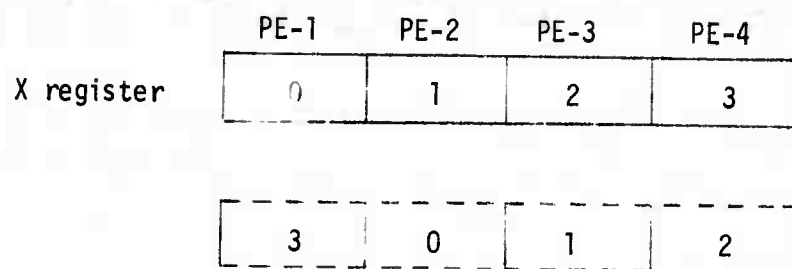
Figure 59a.

Figure 59b.

142
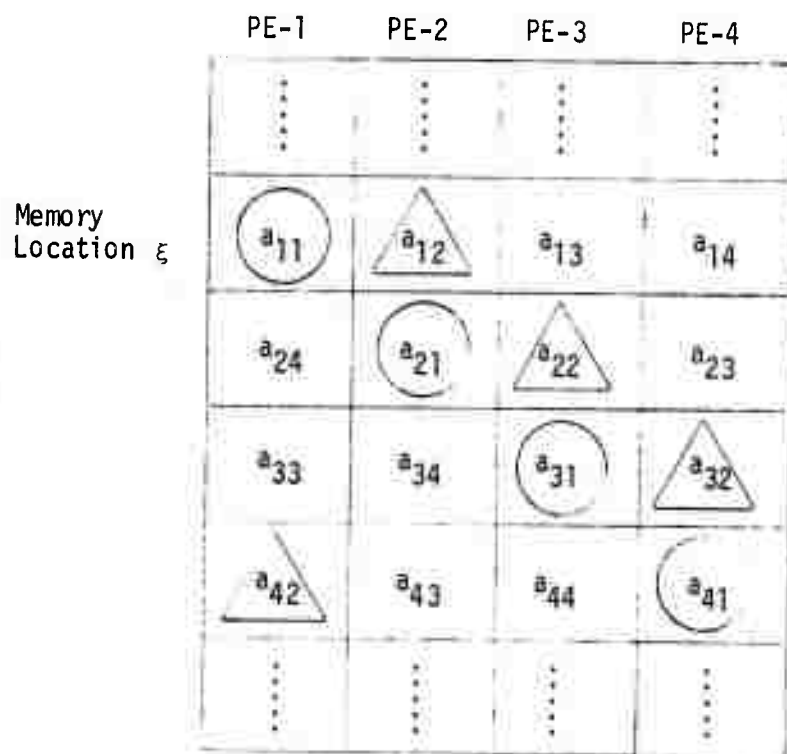
References Cited

1. Slotnick, D. L., et al, The ILLIAC IV Computer, *IEEE Transactions on Computers*, v. C-17, no. 8, p. 746-757, 1968.

2. Denenberg, S. A., An Introductory Description of the ILLIAC IV System, *ILLIAC IV Document No. 225*, University of Illinois, Urbana, Illinois, July 1970.

143

MR. TRULIO: Have you a simple one-D hydro-code, say, which includes an equation of state, running on a machine?

MR. MC INTYRE: No, we haven't. The machine is actually not scheduled to run until sometime this winter. We looked very closely at a two-D hydro-code that the weapons lab runs and found it is fairly straight-forward to adapt to ILLIAC IV as far as the finite-difference scheme is concerned. We only considered a gamma law equation of state. Had we wanted to use a real equation of state it would have been a little more complicated.

CHAIRMAN SIMMONS: Are there other comments or questions? There is another session on the ILLIAC planned for tomorrow, so we might hold off extensive discussion.

MR. GRINE: What is going to be the procedure by which non-University of Illinois people use this?

MR. MC INTYRE: Probably the best procedure is to talk to your ARPA sponsor and he can arrange for time on the machine.

Preceding page blank

SUMMARY OF JUNE 8 SESSION

Gene Simmons
Massachusetts Institute of Technology

We have a few minutes left. It might be useful if I summarized in a very few short statements the conclusions that I have reached from today's discussion. They may not be the same as you have reached.

I knew beforehand that rock mechanics was very complicated, and nothing has been said today that changed my views. I guess every time I talk to people who make code calculations I realize again that there are extremely simplifying assumptions they use, which seem to be necessary whenever we try to take the real data we gather from simple lab measurements and use them in the interpretation of field data.

I am still confused--I guess "confused" is the right word--on how to properly extrapolate data to field situations. I think we touched on several aspects of that problem today, though it was never quite stated in exactly those words. Several of the speakers emphasized the need for in situ measurements. I certainly concur. There have been only a few measurements in situ that have been reported in the literature with corresponding laboratory measurements.

I rather like Hank Cooper's suggestion of a source book. That follows somewhat along one of the suggestions John Handin implied, namely, that there have been a lot of data used from the literature in codes without discrimination. It might be very useful for a few people to call it the most reliable set of data for our purposes. I guess if someone does do this, selection of the most reliable data will depend very strongly upon his own prejudices. The matter of reliability always depends on a subjective evaluation.

I have the feeling that the major question with code calculations is still about what it was two and half years ago; primarily the uncertainties in the input data.

Finally, I am certainly impressed with the potential of the ILLIAC IV. I am a little floored in thinking about having to learn a new language. But the potential for a few of us with the ILLIAC IV certainly is impressive.

RECENT PROGRESS IN THE STUDY OF DYNAMIC ROCK PROPERTIES PERTINENT TO PREDICTING SEISMIC COUPLING*

Thomas J. Ahrens
Seismological Laboratory
California Institute of Technology

Abstract

Recent progress is reviewed in the study of the dynamic yielding of porous and nonporous rocks, the effect of water on the equation of state, and very high pressure equations of state. New Hugoniot data, both above and below the Hugoniot elastic limit, are compared with hydrostatic compression data. This comparison indicates that appreciable stress differences, comparable to those existing at the dynamic yield point, are supportable by rocks, such as sandstone, limestone, and granite above the dynamic yield point. Quasi-static failure tests provide data that closely satisfy a Prager-Drucker-type yield surface. One-dimensional stress tests for a series of porous and nonporous rocks indicate sensitivity of fracture stress and dynamic (Young's) modulus on strain rate. Water is found to affect the dynamic flow and resulting stress waves from underground explosions because of steam formation and as a result of its effect on the phenomenon of block sliding. The latter mechanism is believed to represent the dominant process limiting shear stresses at late times in the flow around explosions in competent rocks. At high shock pressures a series of shock-induced phase transitions involving changes of from 10 to 60 percent in density takes place in silicate minerals. These account for most of the compression that occurs in the first megabar of pressure in solid rocks. As a result of these transitions the release adiabats upon initial unloading from Hugoniot states are considerably steeper in the stress-density plane than the corresponding Rayleigh lines and give rise to appreciable intrinsic shock attenuation.

Introduction

Since the previous review of equation-of-state data relevant to the VELA Uniform Point Source Program (Simmons, 1968)**, available data obtained by both static and dynamic techniques has markedly increased. New and important results have been reported in the study of

---

the effect of irreversible densification--arising from initial porosity, or phase changes, or both. The effect of water on the equation of state of rocks and equations of state of mixtures, in general, are other areas in which important advances have occurred. New data and theories have been reported relating to the static and dynamic failure criteria for rocks. The effect of strain rate on yielding, as well as the importance of the dilatancy phenomenon on the behavior of rocks under the dynamic action of stress waves have only recently been recognized.

Much of the research that has been carried out in these areas has been motivated by the need for providing a complete mechanical and thermodynamic description, i.e., the constitutive relation or equation of state, for a variety of earth materials. A complete equation of state should relate stress (or pressure), strain (or density), and one or more thermodynamic variables, such as internal energy (or temperature) for a specific material. In addition, a complete knowledge of the mechanical yielding conditions and the appropriate post-yield rheological behavior is needed.

This knowledge of the equation of state of earth materials is necessary for the calculation of intense stress-wave propagation and seismic coupling resulting from explosions or from impact on or within the media of the earth's crust. Because of the wide range of dynamic pressures induced within the vicinity of a nuclear explosion, the equations of state that are of interest will describe the response of a medium over a range of conditions from thousands of kilobars and tens of thousands of degrees down to a fraction of a bar at ambient temperature. The latter conditions, of course, correspond to large distances from the disturbance. The spatial relation between the elastic, dynamic yielding, and hydrodynamic zone around an underground explosion is indicated in Figure 60.

At all distances from the explosion, as each point in the medium is encompassed by the outgoing stress or shock wave, it achieves a stress state which usually corresponds to the maximum value it will experience upon passage of the stress wave. Initially, this compression is one-dimensional and corresponds to a state along the Hugoniot of the material. At later times, successive rarefaction states will be achieved which lie along the release isentrope of the material centered at the initial Hugoniot state. Upon stress release, radial flow also takes place and the strain is no longer one-dimensional.

The thermodynamic equation of state of the medium will be of importance to numerical calculations of stress-wave propagation when the stress levels are significantly above those where the dynamic yielding condition is achieved, in the case of rock, and above the stress levels at which a complete compaction has occurred in the case of initially distended materials. These stress levels correspond to several kilobars, in the case of weak and/or porous rocks, and perhaps one-hundred or more kilobars in the case of well indurated rock materials.

148

PASSIVE ELASTIC
WAVE PROPAGATION

ELASTIC
RELAXATION ZONE

{TENSILE FAILURE
{RADIAL CRACKING
{HUGONIOT ELASTIC
{LIMIT — MAX. STRESS
{CRUSHED ZONE
{(HYDRODYNAMICS)
GAS CAVITY

PLASTICALLY
DEFORMED ZONE
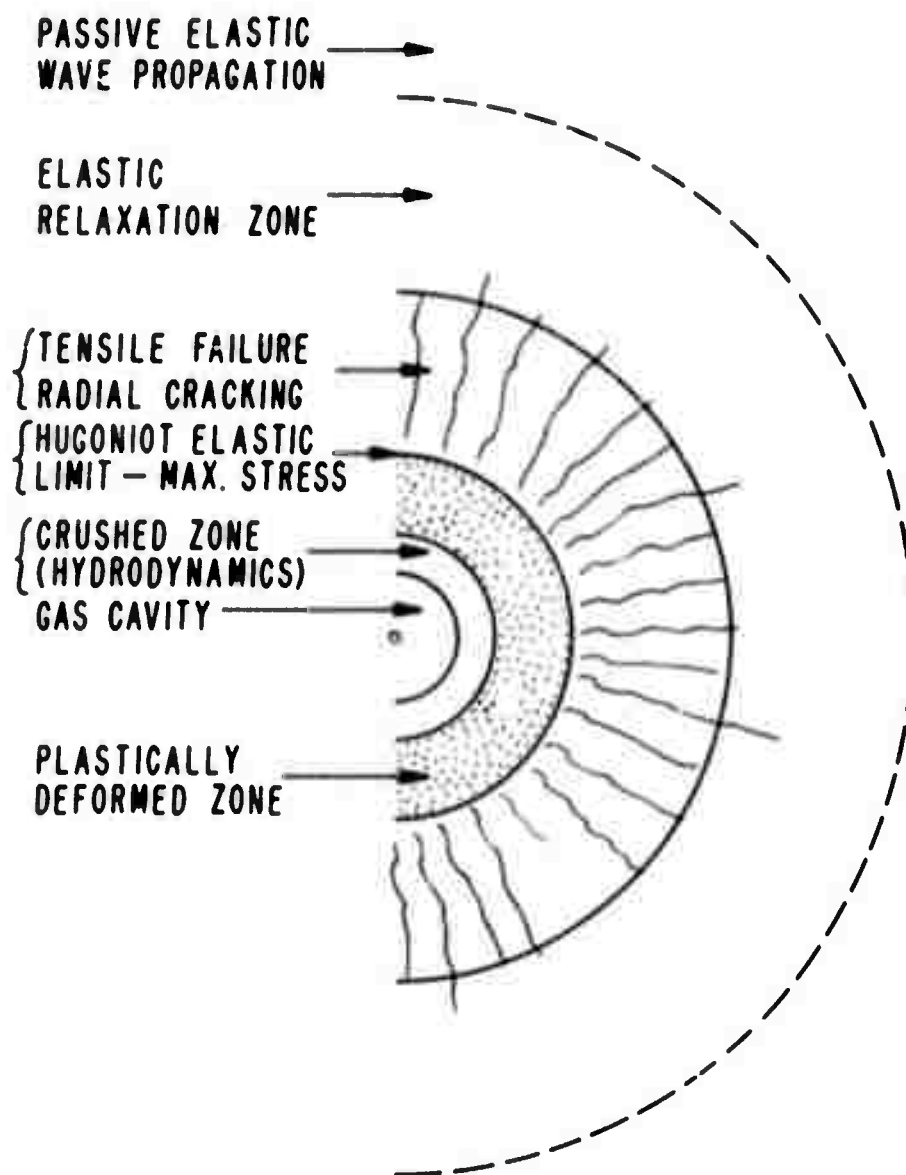
Figure 60. Zones Around an Underground Explosion. Dynamic yielding regime extends from outer boundary of hydrodynamic zone to the limit of radial cracking. Existing tectonic stresses are relieved out to the elastic relaxation zone. In passive elastic wave propagation zone, and beyond, ambient tectonic stress field remains unaffected by explosion. (Modified from C. Archambeau, 1971, private communication).

149

Below these stress levels, within the so-called dynamic-yielding regime, the mechanical effects accompanying dynamic-yielding phenomena dominate over the thermodynamic effects in determining the equation of state of the medium.

In the dynamic-yielding regime, as continual flow takes place during the release process, the tangential stresses will decrease more rapidly than the radial stresses and a state will be achieved which may approach one-dimensional stress conditions. At this point, volume dilation may take place. Hence, the rheological description of the material, particularly in the radial-cracking region (Figure 60), during the latter portions of the flow, must be taken into account. The one-dimensional stress experiments carried out at strain rates significantly lower than those pertinent to compression from shock waves provide important data for describing this regime.

The progress that has been made in the study of dynamic yielding, compaction, and rheology of material behind the shock front is summarized in Sections I and II. The effects of water on the equation of state of rocks, which has recently been the subject of several theoretical and experimental studies, is summarized in Section III. New data pertaining to very high pressure equations of state are summarized in Section IV.

I. Dynamic Yielding and High Pressure Properties of Rocks

The regime of dynamic yielding around an underground explosion extends from an inner radius, within which rock behavior can be closely described in terms of thermodynamic functions and a hydrodynamic-type rheology, to an outer radius at which radial cracking--resulting from the dynamic tensional failure--has ceased. This outer radius, which marks the onset of the elastic or seismic regime, is probably controlled by the initial jointing and cracking geometry within the rock rather than being an intrinsic rock property. These tensile fractures, the outer limits of the dynamic yielding regime, arise from hoop stresses produced by the divergent flow taking place around underground explosion. As indicated in Figure 60, the dynamic-yielding zone is thus divided into the outer (tensional failure) zone and the inner zone of dynamic shear and compressive failure. The stress level in the initial portion of the stress pulse, which suggests the subdivision between the two zones of failure, is the Hugoniot elastic limit. Since in the initial portion of the stress wave, no radial divergence of the flow has occurred, it is convenient to define the Hugoniot elastic limit as: the maximum stress which may be achieved upon the rapid one-dimensional compression without internal rearrangement taking place at the shock front.

150

Measurements of the Hugoniot elastic limits (HEL) in the last two years have been obtained for Hardhat granite (~40 kb) (Figure 61), Cedar City (tonalite) granite (15-20 kb) (Figure 62), Solenhofen limestone (~6 kb) (Figure 63), Pictured Cliffs sandstone (Figure 64), and single crystal halite (Murri and Anderson, 1970). For the granite, sandstone, and limestone hydrostatic compression measurements are also available at sufficiently high stress levels to be comparable to Hugoniot shock states above the Hugoniot elastic limit. Such data are useful in formulating a rheological model for the rock above the HEL. Petersen et al. (1968) show that in the case of the Pictured Cliffs sandstone, the Hugoniot states above the HEL are offset above the hydrostat by a stress which is 4/3 times the maximum shear stress measured in quasi-static triaxial strength tests. This result (Figure 64) is in accord with simple elastoplastic theory which predicts

$$\sigma_{off} = \frac{4}{3} \tau_{max} \tag{1}$$

Also σ_m, the mean stress, is given by

$$\sigma_m = (\sigma_x + 2\sigma_y)/3 \tag{2}$$

The lateral stress (parallel to the shock front) for one-dimensional compression is given from elasticity theory by

$$\sigma_y = \nu\sigma_x/(1 - \nu) \tag{3}$$

where ν is the Poisson's ratio. Similarly, the maximum shear stress is

$$\tau_{max} = \frac{1}{2} (\sigma_x - \sigma_y) \tag{4}$$

In the case of Hardhat granite (Figure 61) the calculated stress offset of the Hugoniot above the hydrostat is on the order of 15 kb. This is somewhat less than the 24-kb offset predicted from elastoplastic theory using a Poisson's ratio of 0.22 and a Hugoniot elastic limit of 50 kb (Cherry and Petersen, 1970). This result implies that, at least behind the deformational shock front, the rock can retain the stress differences of ~22 kb, slightly less than the ~36 kb stress difference which is supportable just at the dynamic yield point.

The shock and ultrasonic data for Solenhofen limestone of Jones and Froula (1968) and Peselnick (1962) may be similarly used to predict the value of σ_{off}. Using a 2.57 g/cm^3 density value for the Solenhofen limestone from Peselnick's table of elastic constants versus density, yields a Poisson's ratio of 0.294. With this value, a stress offset of ~2.3 kb is calculated for a Hugoniot elastic limit of 6 kb (Figure 63).

151
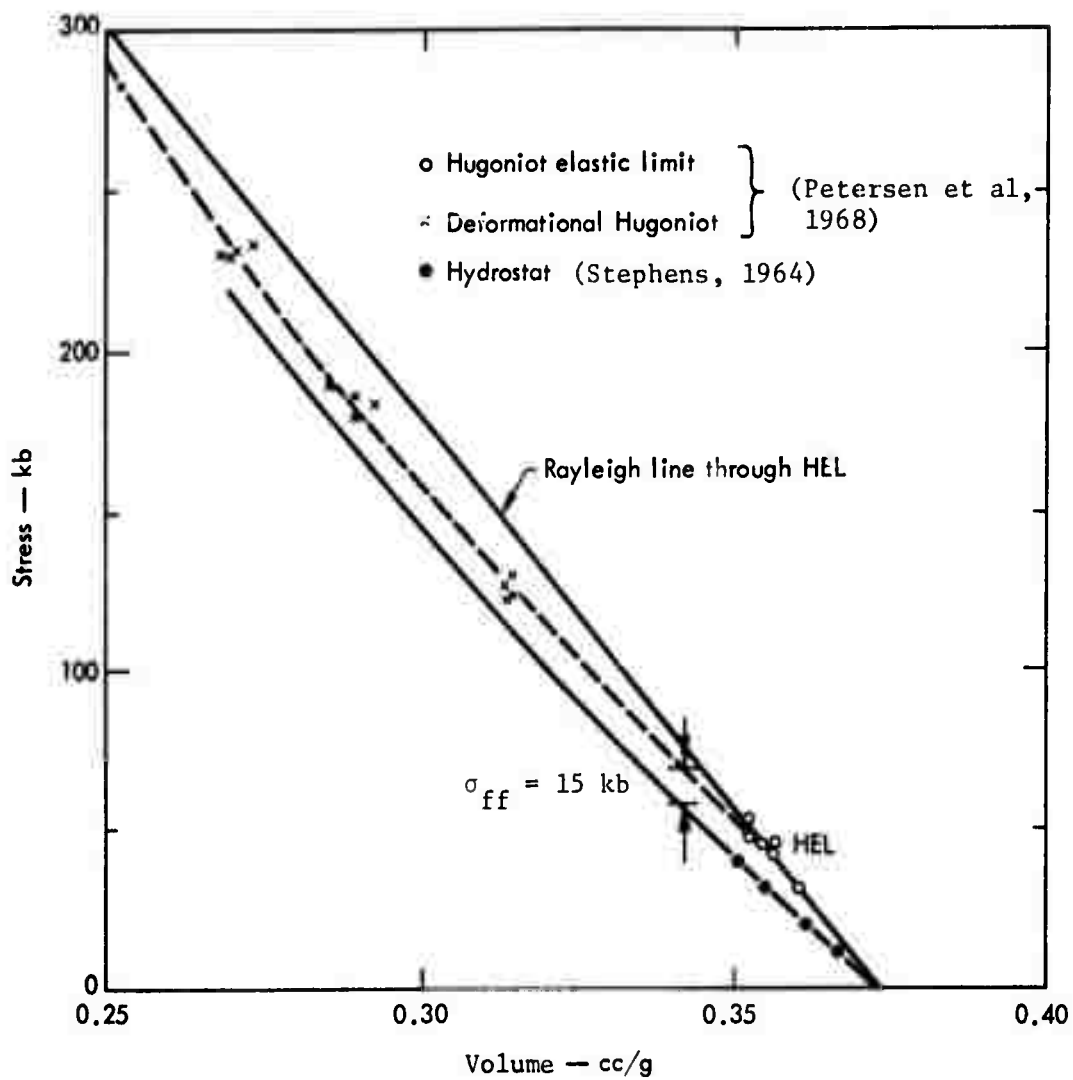
Figure 61. Hugoniot and Hydrostatic Compression Data for Hardhat Granite. (After Cherry and Petersen, 1970).
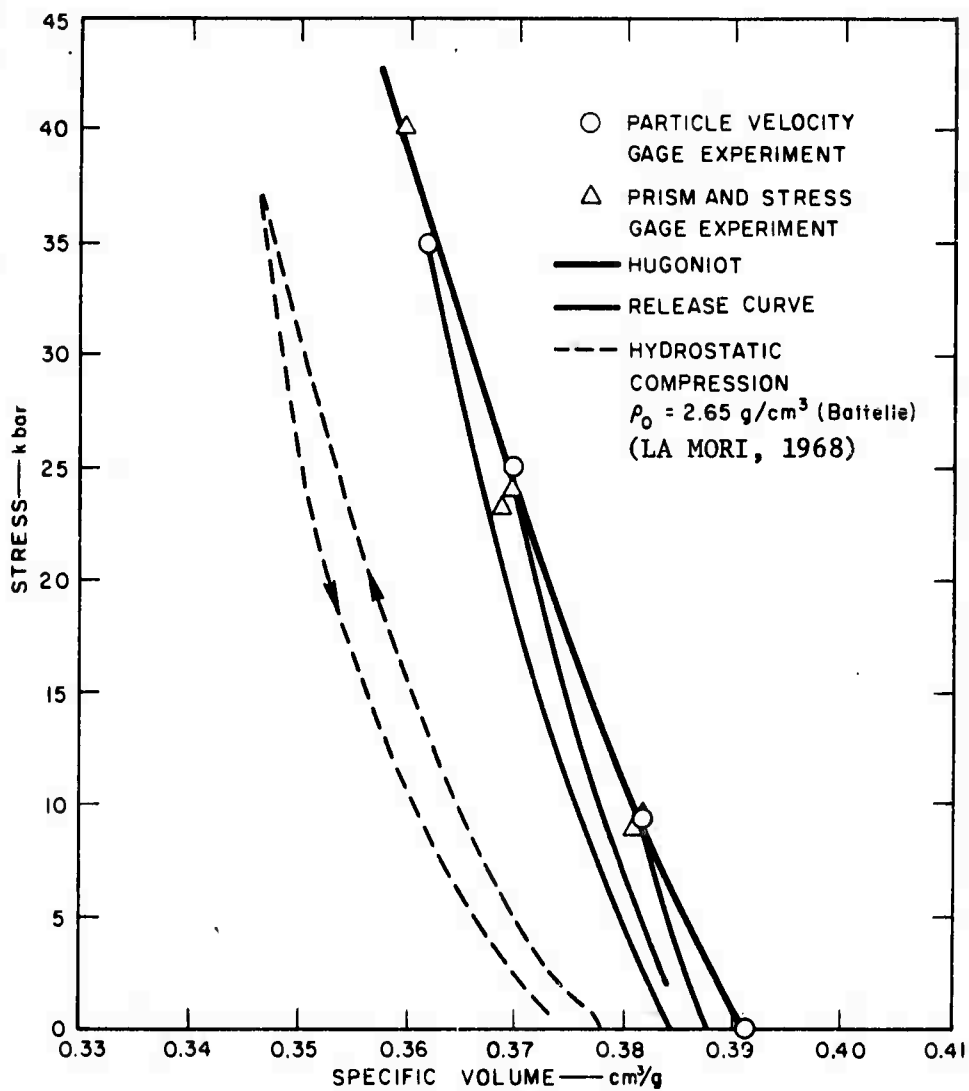
Figure 62. Hugoniot and Release Adiabat Data for Cedar City Granite
(Tonalite). (Modified from Petersen et al., 1968).

153
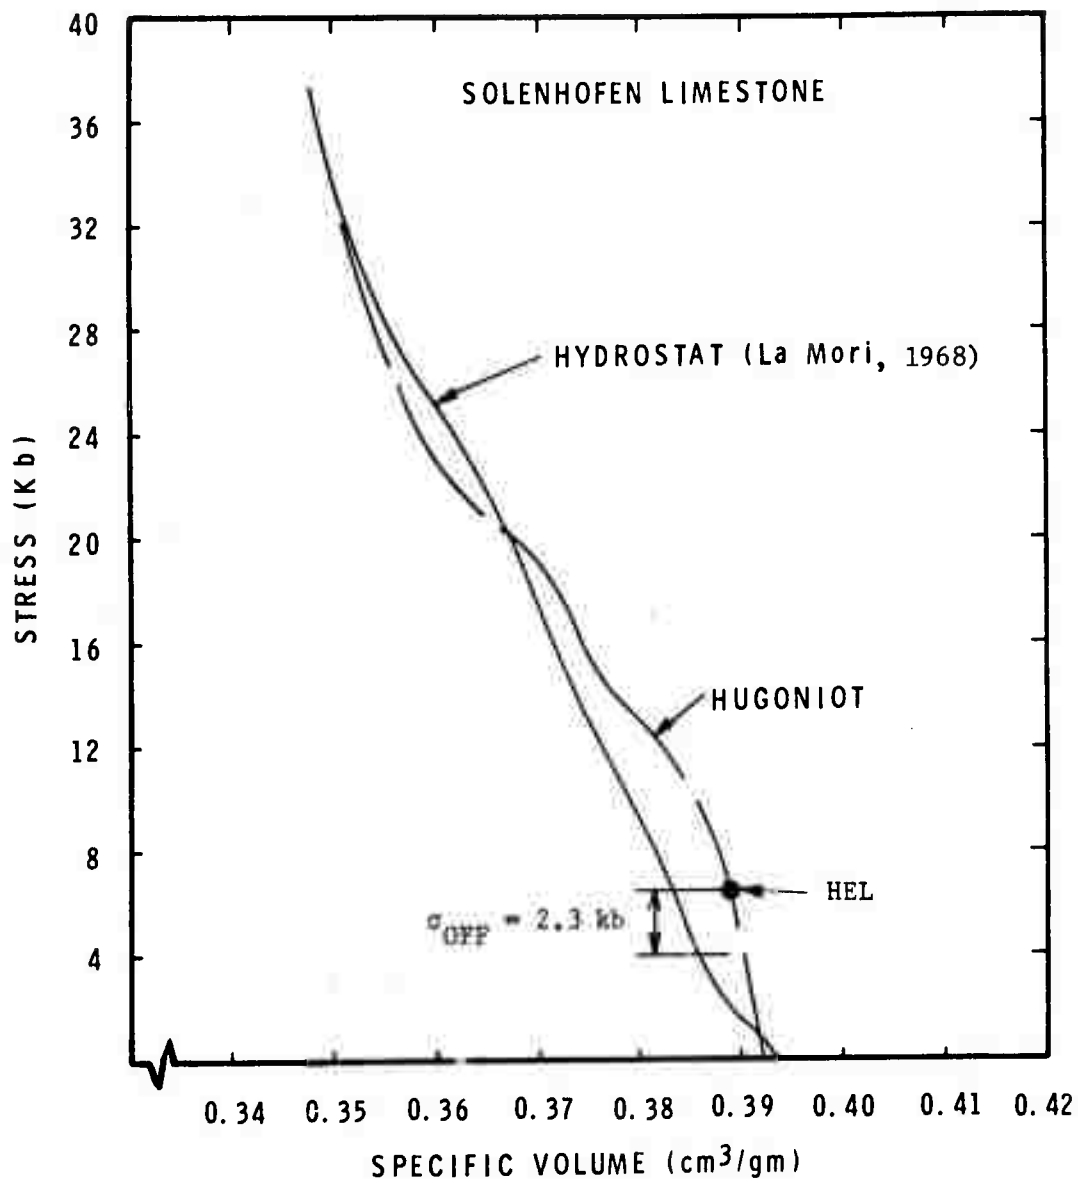
Figure 63. Hugoniot and Hydrostatic Compression Data for Solenhofen Limestone. (After Jones and Froula, 1968).
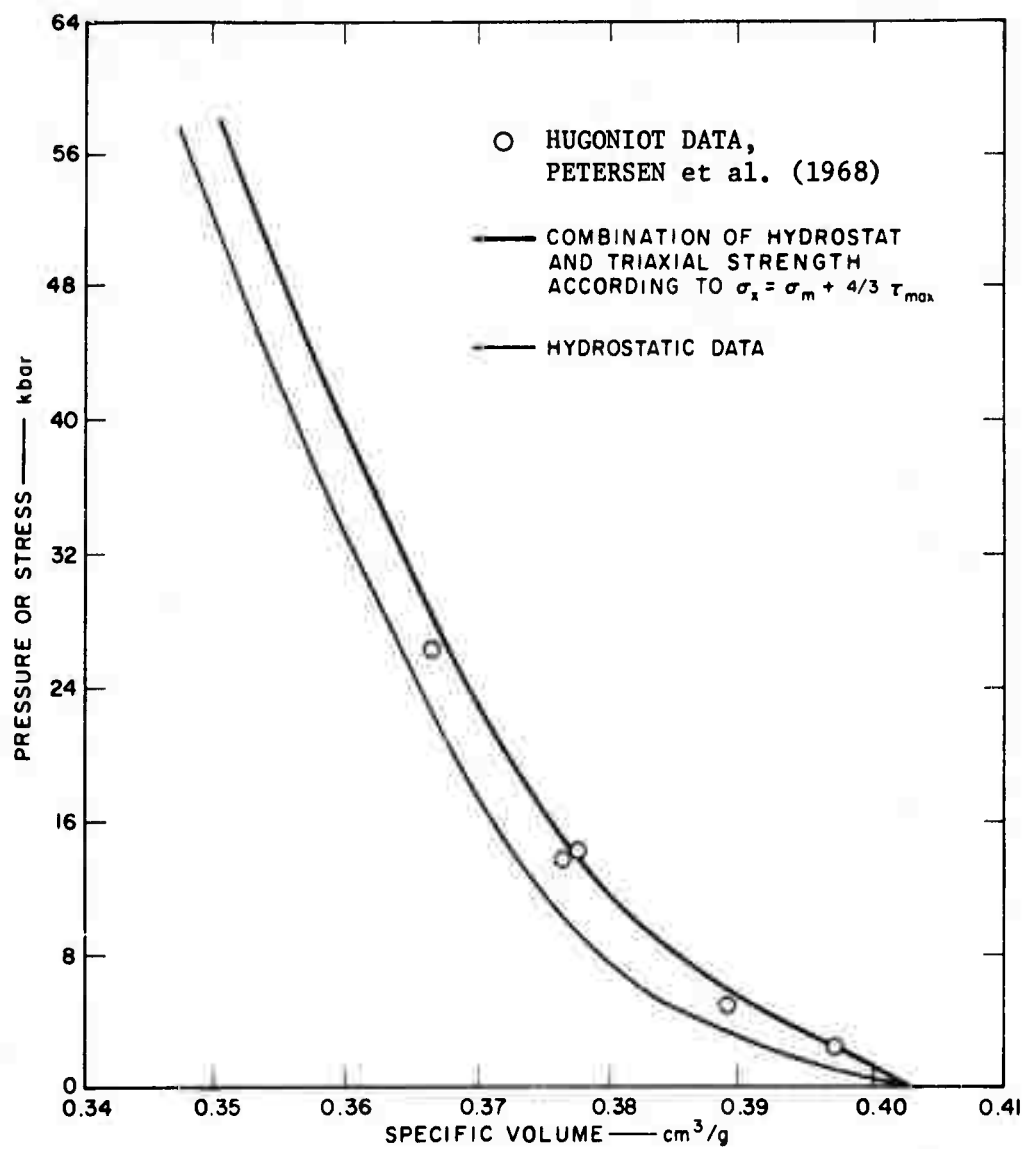
Figure 64. Hugoniot and Hydrostatic Compression Data (taken by Stephens et al., 1970) for Pictured Cliffs Sandstone. (After Petersen et al., 1968).

155

Comparison with LaMori's (1968) compression curves on a slightly more porous sample indicates a stress offset of ~5 kb at the Hugoniot elastic limit, which decreases to essentially zero at ~20 kb. At higher shock stresses, the Hugoniot appears to lie slightly <u>below</u> the hydrostat, which may result from the onset of several complex shock transitions. These transitions are probably shear-stress activated, within the ~15 to ~30-kb interval (Ahrens and Gregson, 1964). The dynamic-yielding data for single crystal halite (Murri and Anderson, 1970) has yielded values of the HEL varying from 0.2 to 0.3 kb in the [100] direction to 7 to 8 kb in the [111] direction. This variation in HEL with orientation is explained in terms of activation of specific slip systems upon compressing along different crystallographic axes. These results are significant in that, for the first time, the dynamic-yield mechanism under shock loading of geological material has been explained in terms of a definitive microscopic process.

In addition to the new data for dynamic yielding of rocks, several recent studies have dealt with formulating the yield criterion of rocks under quasi-static and faster strain-rate conditions. The work of Mogi (1967) has shown that the maximum shear stress for quasi-static failure, τ_{max}, can be related to all three principal stresses by a relation of the form

$$\tau_{max} = (\sigma_1 - \sigma_2)/2 = f_1[(\sigma_1 + \sigma_3 + \alpha\sigma_2)/2] \qquad (5)$$

where α is 0.1 to 0.2 for the rocks tested and $\sigma_1 > \sigma_2 > \sigma_3$. The static triaxial strength test data for a sandstone, a limestone, and Westerly granite (Figures 65-68) satisfy this simple modified Tresca model markedly well. The data for simple tension and compression are closely ordered by a function such as suggested by Equation 5. DiMaggio and Sandler (1970) have recently suggested a series of more complex models, including a modified Drucker-Prager model in which the yield surface is given by

$$\sqrt{J_2'} - k + \alpha J_1 \left(1 + \frac{J_1}{2c}\right) = 0 \qquad (6)$$

$$\text{for } J_1 + c > 0$$

where k, c, and α are constants and the stress invariants are given by

$$J_1 = \sigma_1 + \sigma_2 + \sigma_3 \qquad (7)$$

$$J_2' = (\sigma_1 - \sigma_2)^2 + (\sigma_2 - \sigma_3)^2 + (\sigma_3 - \sigma_1)^2. \qquad (8)$$

A yield equation, of the form of Equation 6, appears to give a close
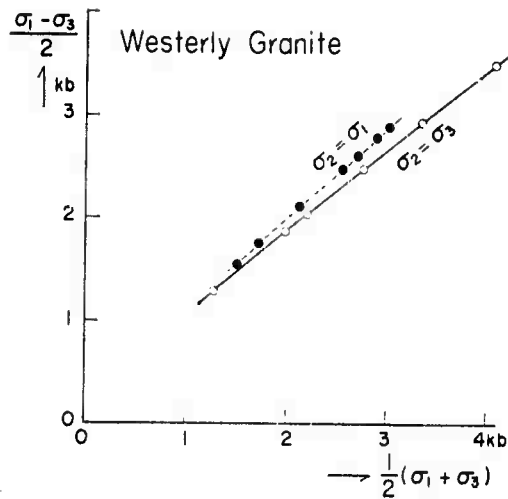
Figure 65. Maximum Shear Stress Versus $(\sigma_1 + \sigma_2)/2$ for Failure of
Westerly Granite. Open circles, compression; solid circles,
extension. (After Mogi, 1967).
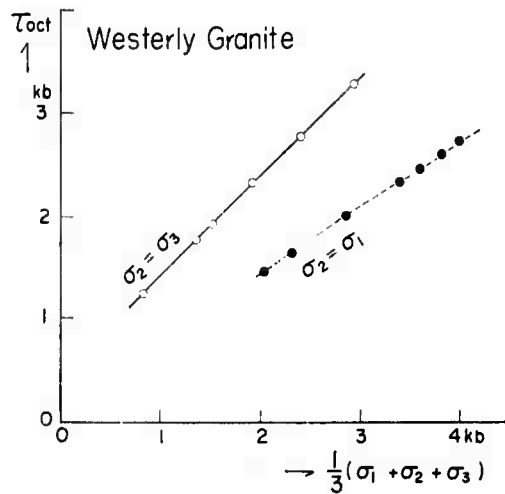
Figure 66. Octahedral Shear Stress τ_{oct} Versus Mean Normal Stress for
Failure of Westerly Granite. Open circles, compression;
solid circles, extension. (After Mogi, 1967).

157
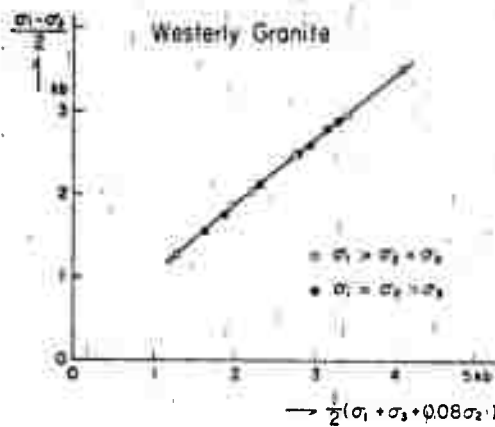
Figure 67.   Maximum Shear Stress $(\sigma_1 - \sigma_3)/2$ Versus $(\sigma_1 + \sigma_3 + 0.08\sigma_2)/2$ for Failure of Westerly Granite.   Open circles, compression; solid circles, extension.   (After Mogi, 1967).
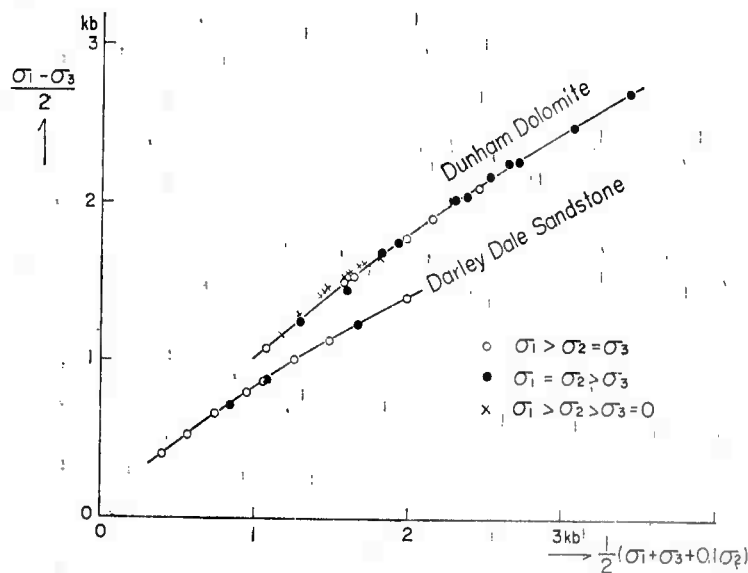
Figure 68.   Maximum Shear Stress $(\sigma_1 - \sigma_2)/2$ Versus $(\sigma_1 + \sigma_3 + 0.1\sigma_2)/2$ for Failure of Dunham Dolomite and Darley Dale Sandstone. Open circles, compression; solid circles, extension; cross, biaxial compression.   (After Mogi, 1967).

158

account of the failure surfaces for a variety of loading paths for rocks as well as some soils. This is demonstrated by results obtained in a series of quasi-static failure tests performed by Swanson (1969) (Figures 69 and 70) which depict the failure envelopes for Cedar City and Westerly granite. In these tests, failure was obtained via several different stress loading paths. Also of interest are results for a competent rock such as Cedar City granite (Figure 71) compressed under one-dimensional strain conditions (the same conditions as for shock compression to states below the Hugcniot elastic limit). For an axial stress of up to ~12 kb, the rock does not fail but demonstrates a clear hysteretic stress-strain curve (Figure 71). Recently Brace (1970) has extended these results to over ~20 kb maximum principal stress. Hugoniot data obtained by Jones and Froula (1968) and Froula (1968) for Westerly granite and anorthosite and for tonalite (Cedar City granite) by Petersen, Murri, and Gates (1969) below the Hugoniot elastic limit, show a similar, essentially linear, stress-strain curve. The results of both Swanson and Brace indicate that only minor hystereses occur in rocks with initial porosities less than 2 percent upon static one-dimensional compression. For more porous rocks, Brace's data demonstrate that irreversible compaction takes place in one-dimensional compression and this compaction is, in some cases, time dependent.

The high levels of dynamic yield strength for granite appear to be applicable to the description of explosions in this material only for short times after the stress wave has enveloped a given volume element. Recently, McKay and Godfrey (1969) have carried out some numerical experiments in which they match observed pressure profiles obtained from small-scale experiments in a series of rock materials. They find that although the instantaneous dynamic yield strength is quite high, as radial expansion of material occurs, sliding of rock blocks within a given shell around the source takes place. They develop a model which assumes that a block sliding mechanism will dominate the rheological properties at late times and give a relatively low overall strength to the medium.

In the last two years, a series of one-dimensional stress experiments, under varying strain rates, have been carried out on different geologic materials. Because of the radially diverging flow around underground explosions, material in the radial-cracking region will enter into a regime in which the stress, rather than the strain, becomes nearly one-dimensional, and on further deformation volume dilatation will probably occur at late times in the flow. Green and Perkins (1968) found that under rapid one-dimensional stress conditions, the dynamic modulus was markedly dependent on strain rate in both a porous rock, i.e., volcanic tuff (Figure 72), and a nonporous rock, such as Westerly granite (Figure 73). In contrast, Solenhofen limestone, which has slight porosity, shows no such dependence on strain rate (Figure 74). A slight increase in yield strength with strain rate is also observed (Figure 75) for this material. It is not yet clear which factors result in strain rate dependent behavior.
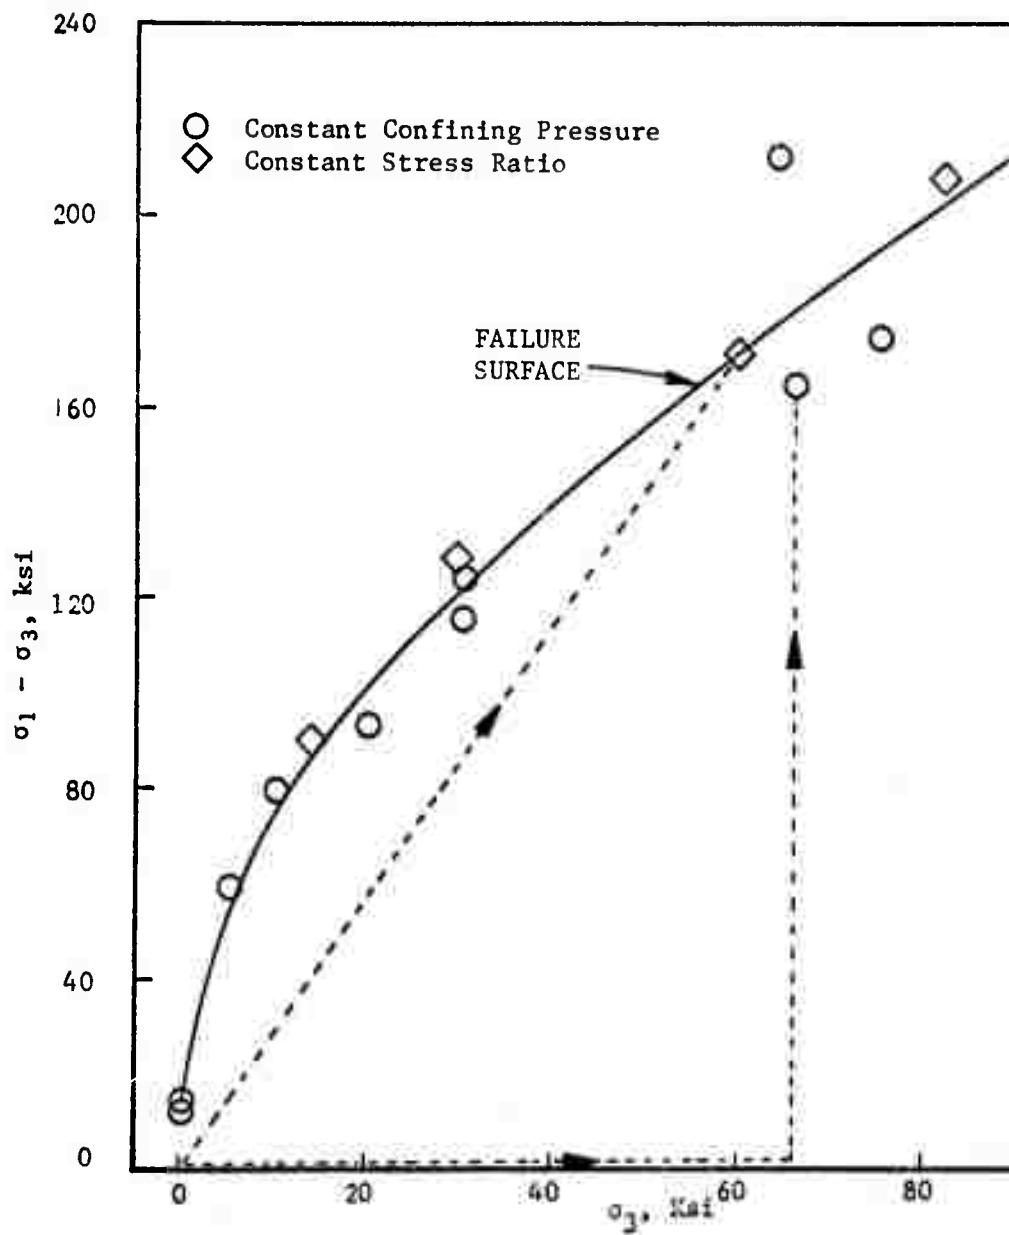
Figure 69. Failure Data for Cedar City Granite Under Constant Confining
Pressure and for Constant Stress Ratio Loading. (After
Swanson, 1969).
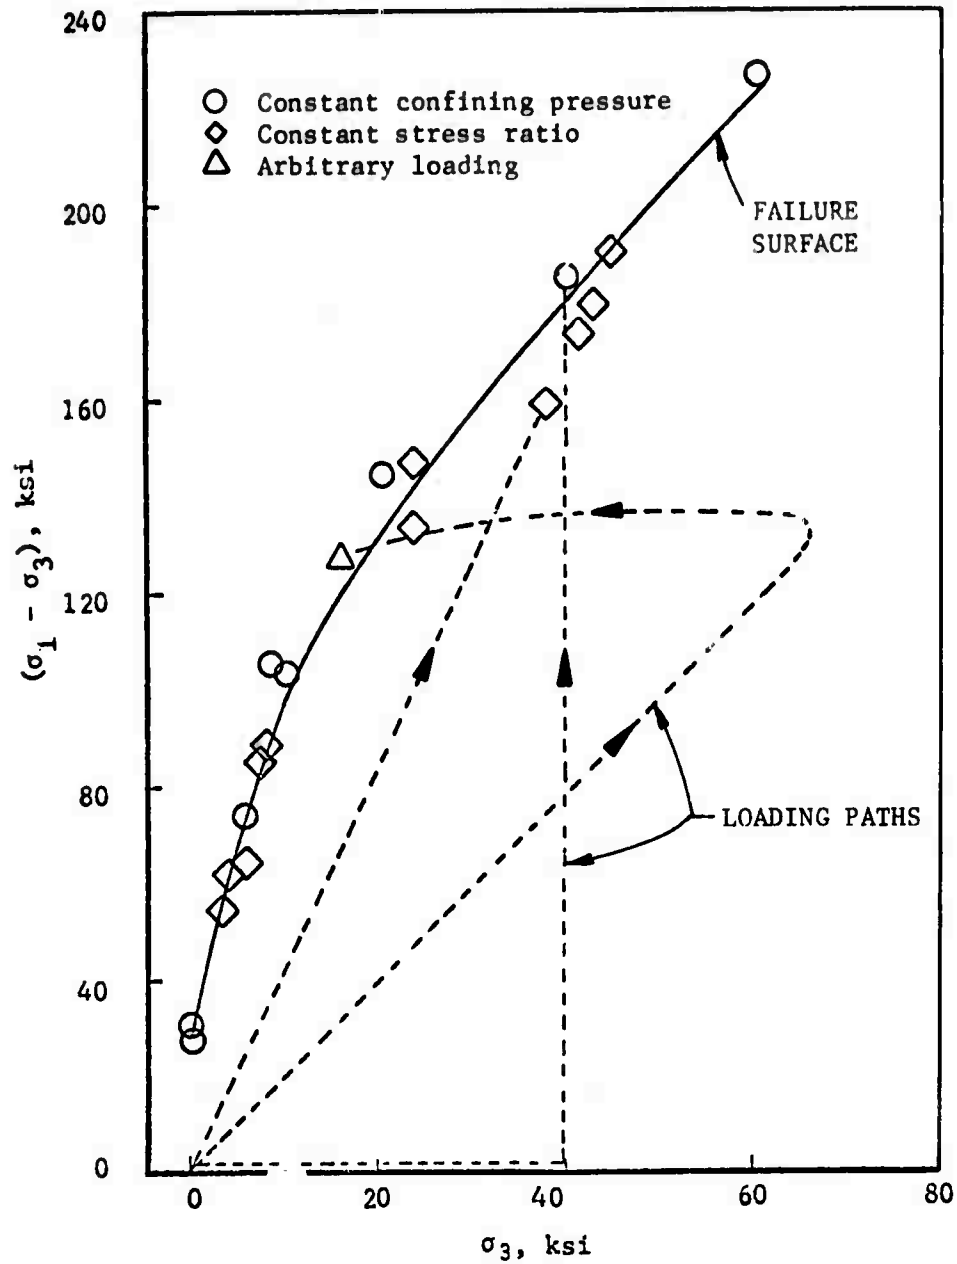
Figure 70. Failure Envelope for Westerly Granite Showing Independence of Loading Path. (After Swanson, 1969).
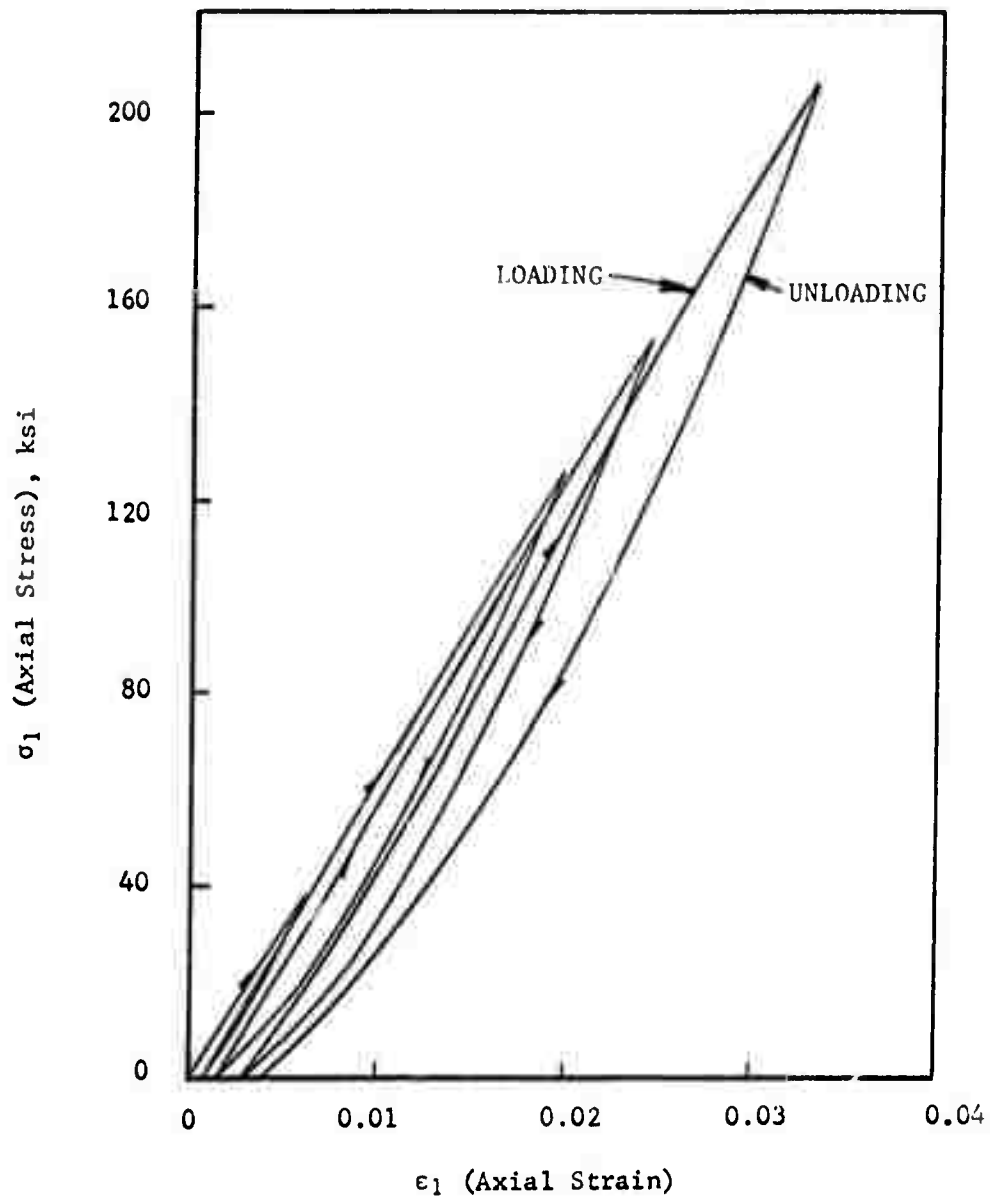
Figure 71. Axial Stress Versus Axial Strain Curves for Cedar City Granite, Tested Under One-Dimensional Strain. (After Swanson, 1969).

162
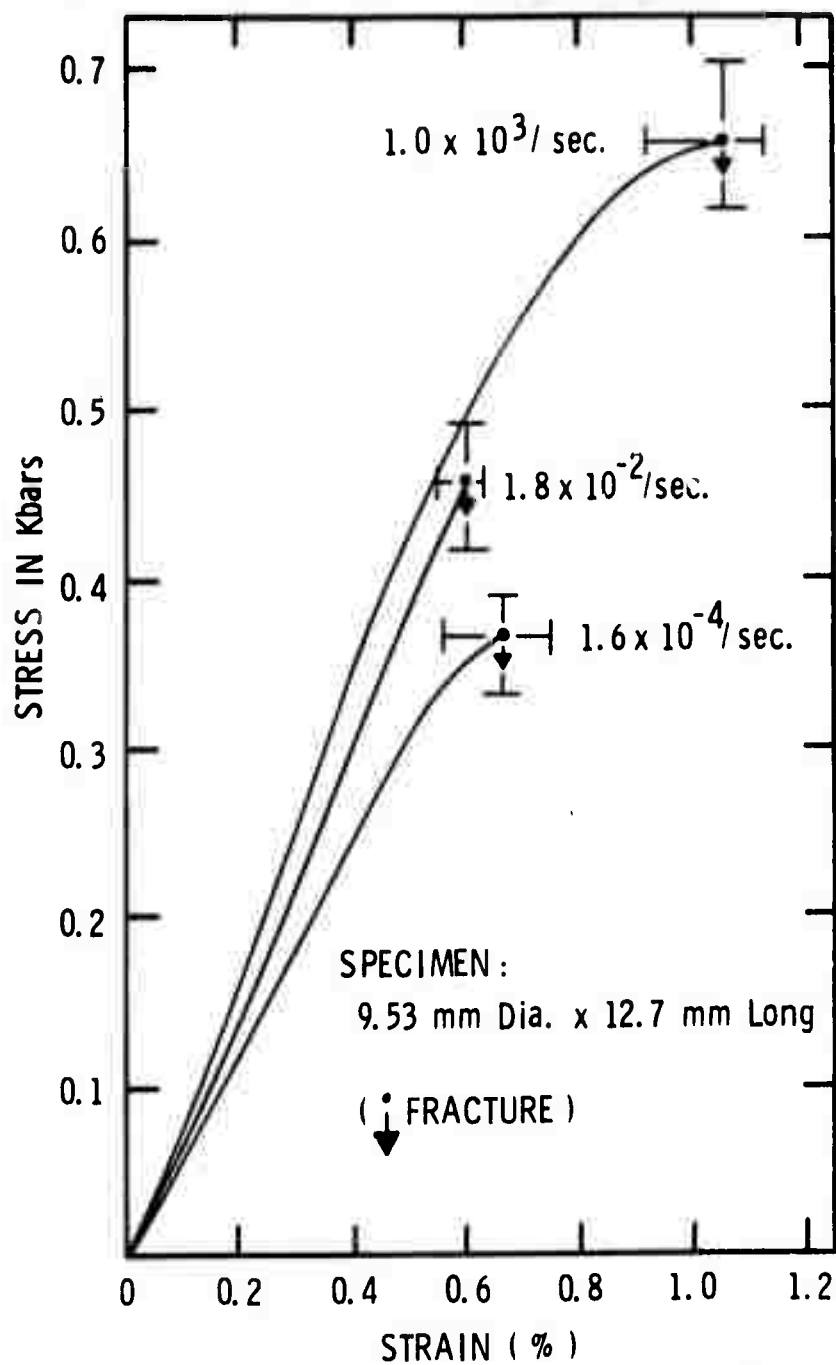
Figure 72. Stress Versus Strain at Various Strain Rates for Volcanic Tuff (dry), Tested Under One-Dimensional Compressional Stress. (After Green and Perkins, 1968).

163
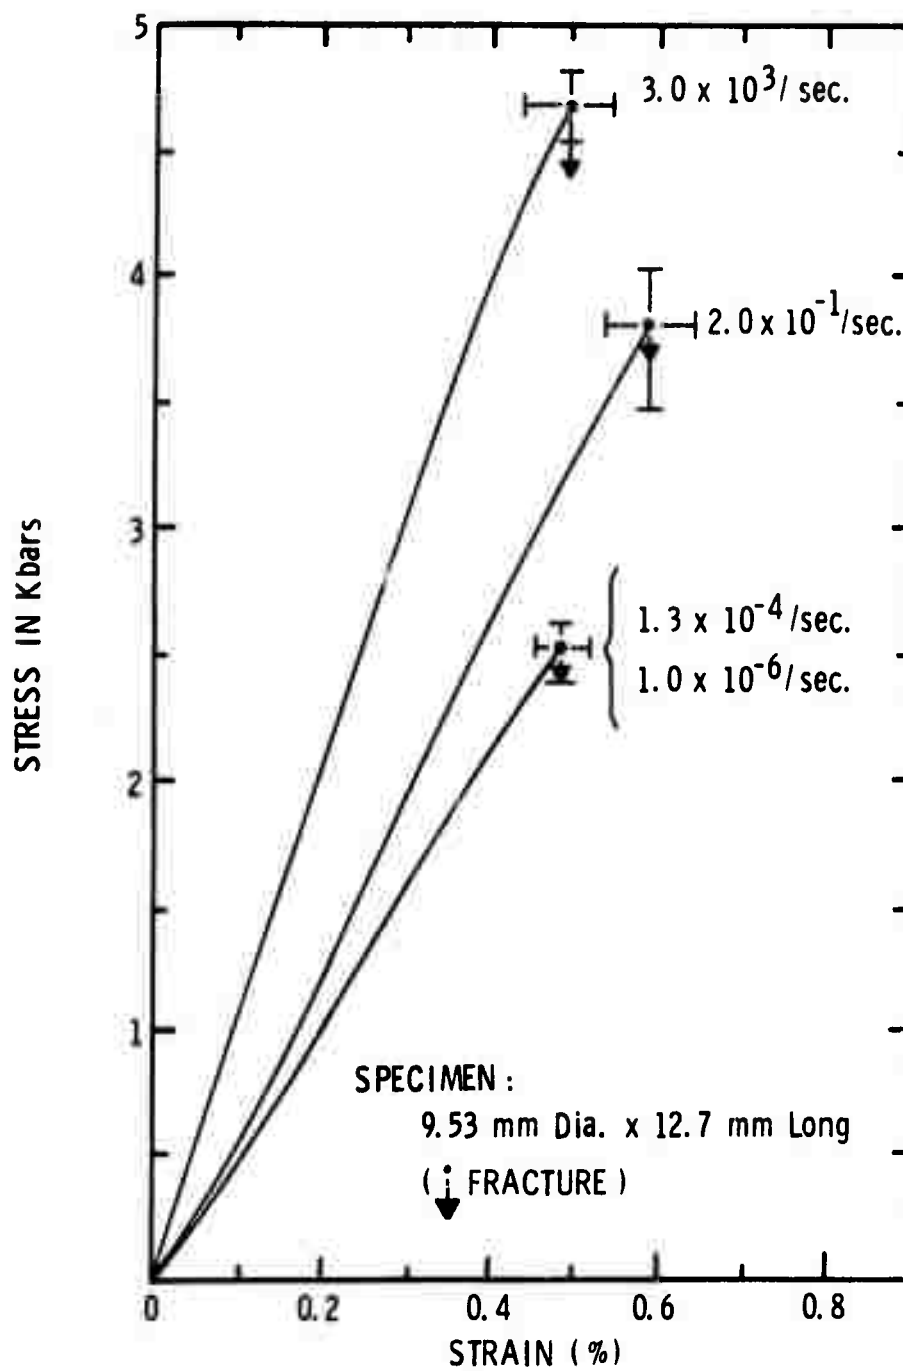
Figure 73. Stress Versus Strain at Various Strain Rates for Westerly
Granite, Tested Under One-Dimensional Compressional Stress.
(After Green and Perkins, 1968).
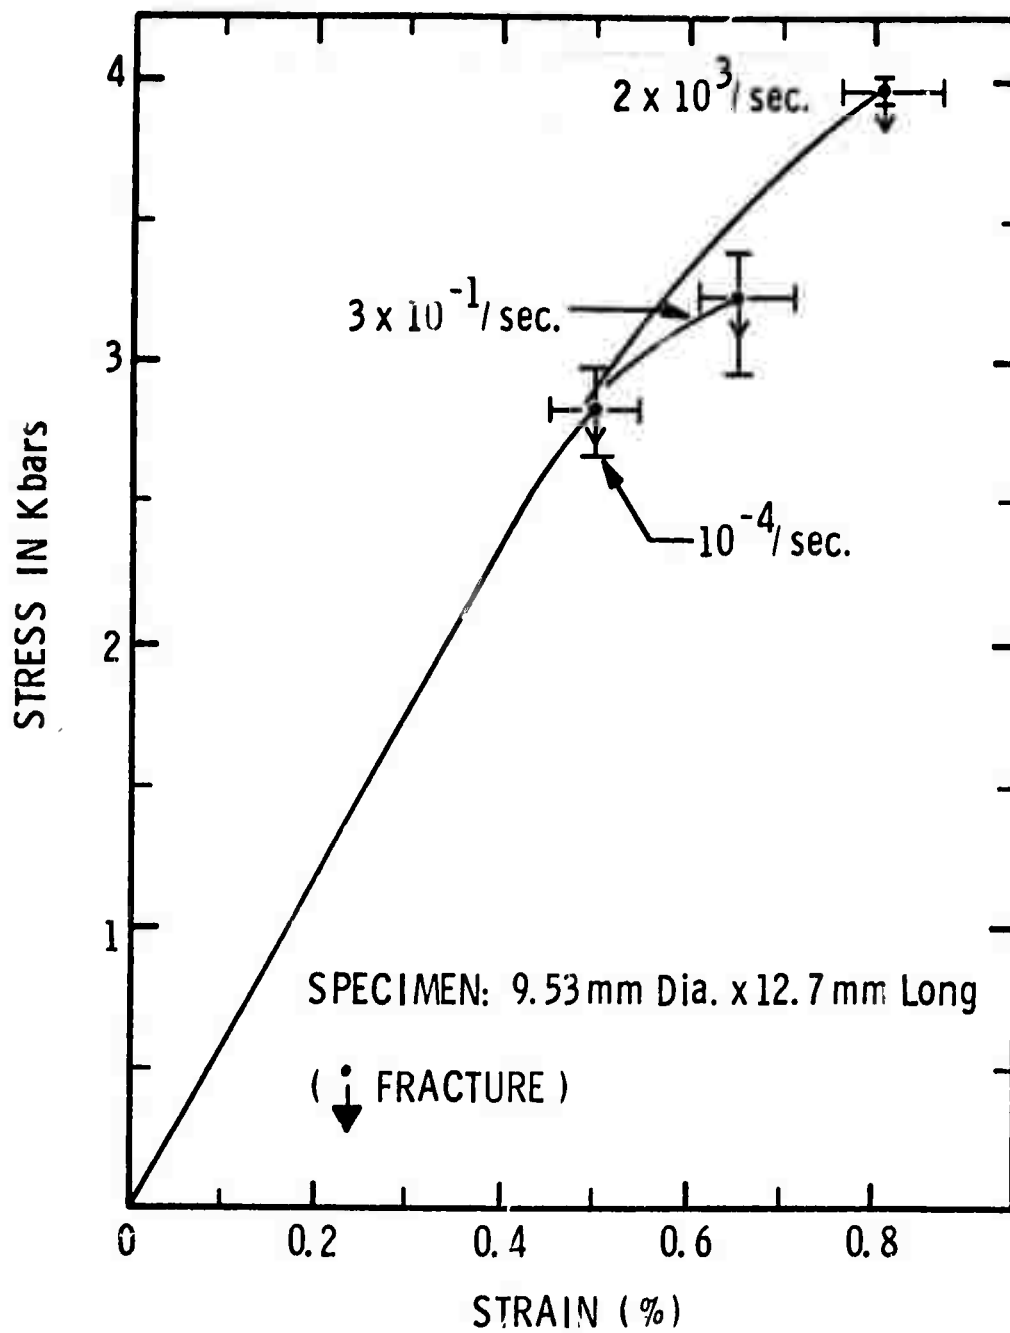
Figure 74.   Stress Versus Strain for Various Strain Rates for Solenhofen
Limestone, Tested Under One-Dimensional Compressional Stress.
(After Green and Perkins, 1968).
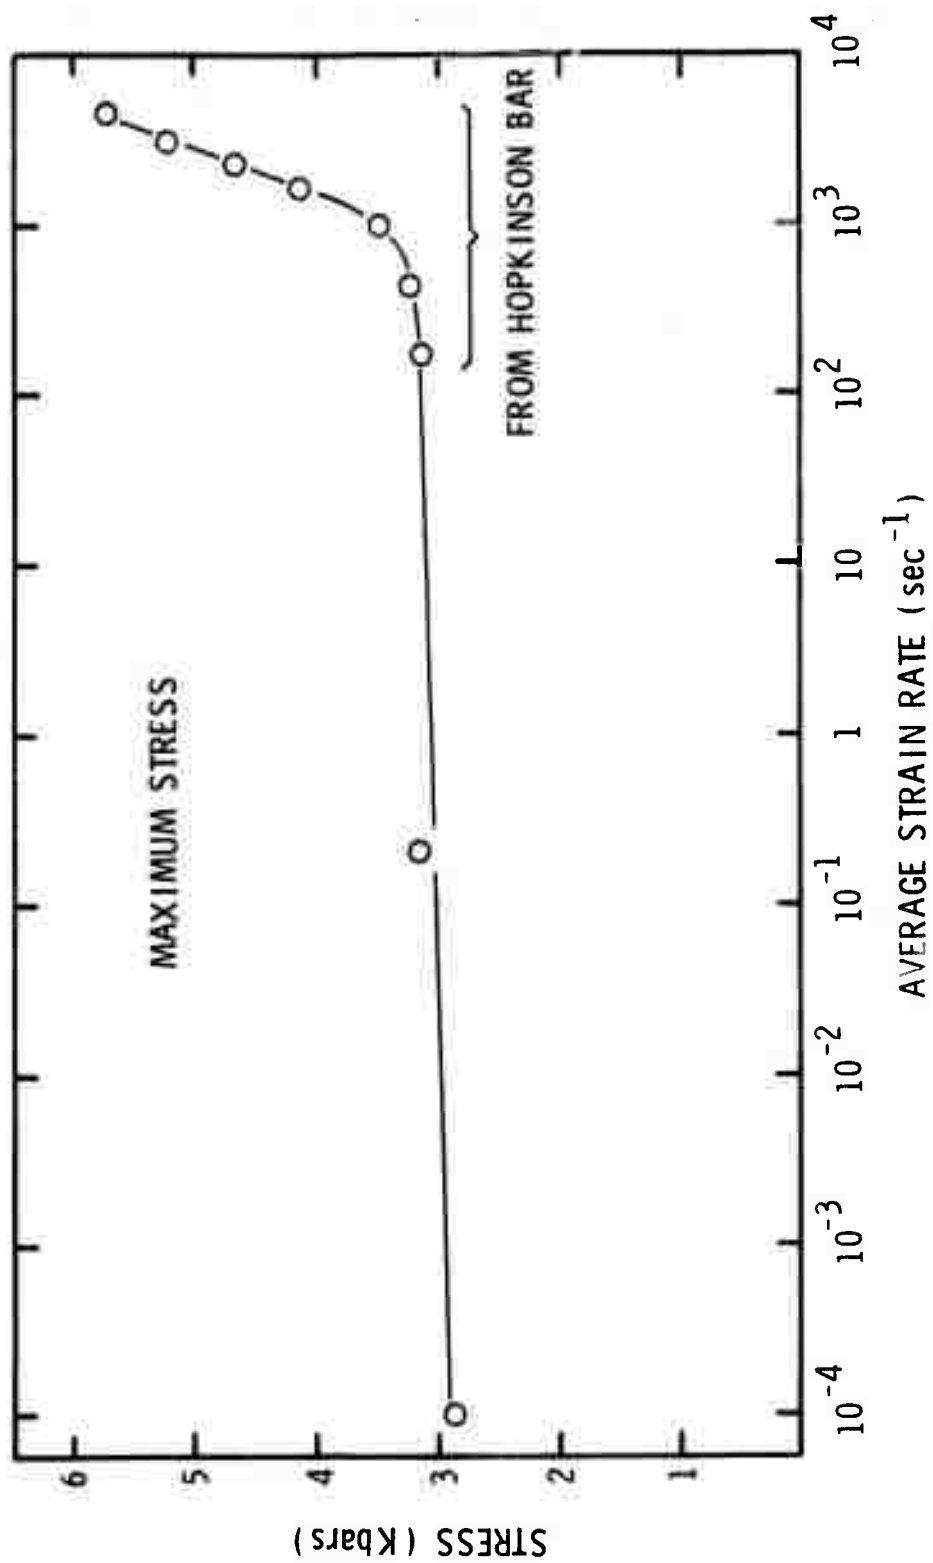
Figure 75. Fracture Stress Versus Log Strain Rate for Solenhofen Limestone Under One-Dimensional Stress. (After Green and Perkins, 1968).

166

The normal <u>decrease</u> in volume followed by a volume <u>increase</u>
(dilatancy) with increasing one-dimensional stress is shown for the
Cedar City granite in Figure 76. Similar results have been observed in
many other rocks upon compression, when these have been instrumented
with strain gages in two directions. The importance of this phenomenon
in describing the later portions of the flow from underground explosions
has only recently been discussed in mathematical terms (McKay et al.,
1970).

II. Porous Materials

The study of the irreversible dynamic compaction of porous
materials by shock waves has occupied several researchers in different
laboratories during the last few years. Among porous materials, volcanic
tuff has been most frequently studied because of the large number of ex-
plosions that have been performed in this medium. One rather surprising
result has been the good agreement in the Hugoniots of wet and dry tuff
at high pressure. The variation between different samples appears to
put at least as much scatter in the Hugoniot data as the effect produced
by changing the water content (Figure 77). The routine application of
the buffer shock reflection method, in which a shock wave in the tuff
is transmitted to a series of low impedance, buffer materials, such that
a rarefaction wave is induced in the sample, and a shock is propagated
to the buffer material, has been applied to dry tuff by Petersen, Murri,
and Gates (1969) and to wet tuff by Rosenberg, Ahrens, and Petersen
(1968). Their release adiabat data are shown in the pressure-particle
velocity plane and pressure-volume planes in Figures 78-81. It is
interesting that not all the permanent compaction which is observed can
be attributed to the irreversible crushing-out of initial porosity.
Pressure-particle velocity and stress-volume release data for fused
quartz (Figures 82-83), which chemically approximates the glassy matrix
of the various NTS tuffs, also behaves in an irreversible manner. The
reason for this is not yet clear, but recent recovery experiments demon-
strate that in all the glassy silicates irreversible densification takes
place as a result of shock compression (Gibbons and Ahrens, 1971).
Release adiabats, which also demonstrate irreversible crushing for a
series of tuffs, with varying initial porosities, have recently been
reported by Lysne (1970) (Figure 84). In this study the novel technique
of measuring release adiabats by observing reverberations of release
waves against a high impedance medium was employed. The technique of
embedding a metallic foil within the rock sample and measuring the
voltage signal, hence the particle-velocity profile, in a sample when
it is shocked in a transverse magnetic field (Dremin et al., 1962) has
provided a new and extremely powerful tool for tracing out a complete
release adiabat of a rock in a single experiment. Representative
results, obtained with this promising technique, for alluvium and
tonalite (Cedar City granite) are shown in Figure 85 and Figure 62.
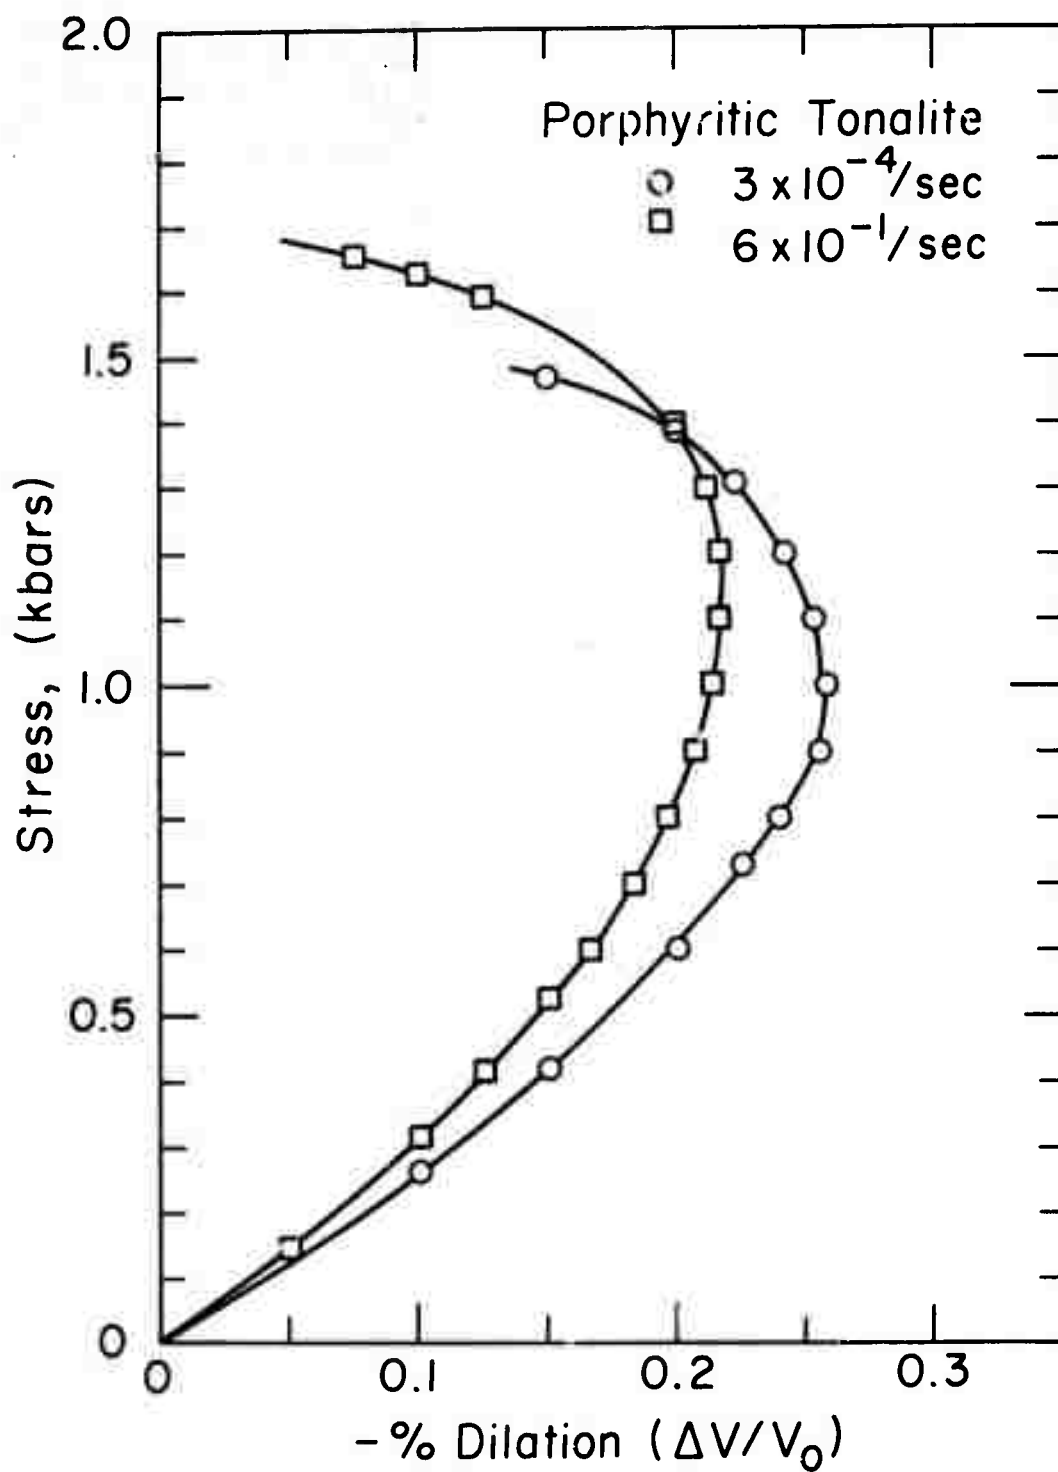
Figure 76. Volume Change as a Function of One-Dimensional Stress for Cedar City Granite (Tonalite). (After Perkins and Green, 1968).
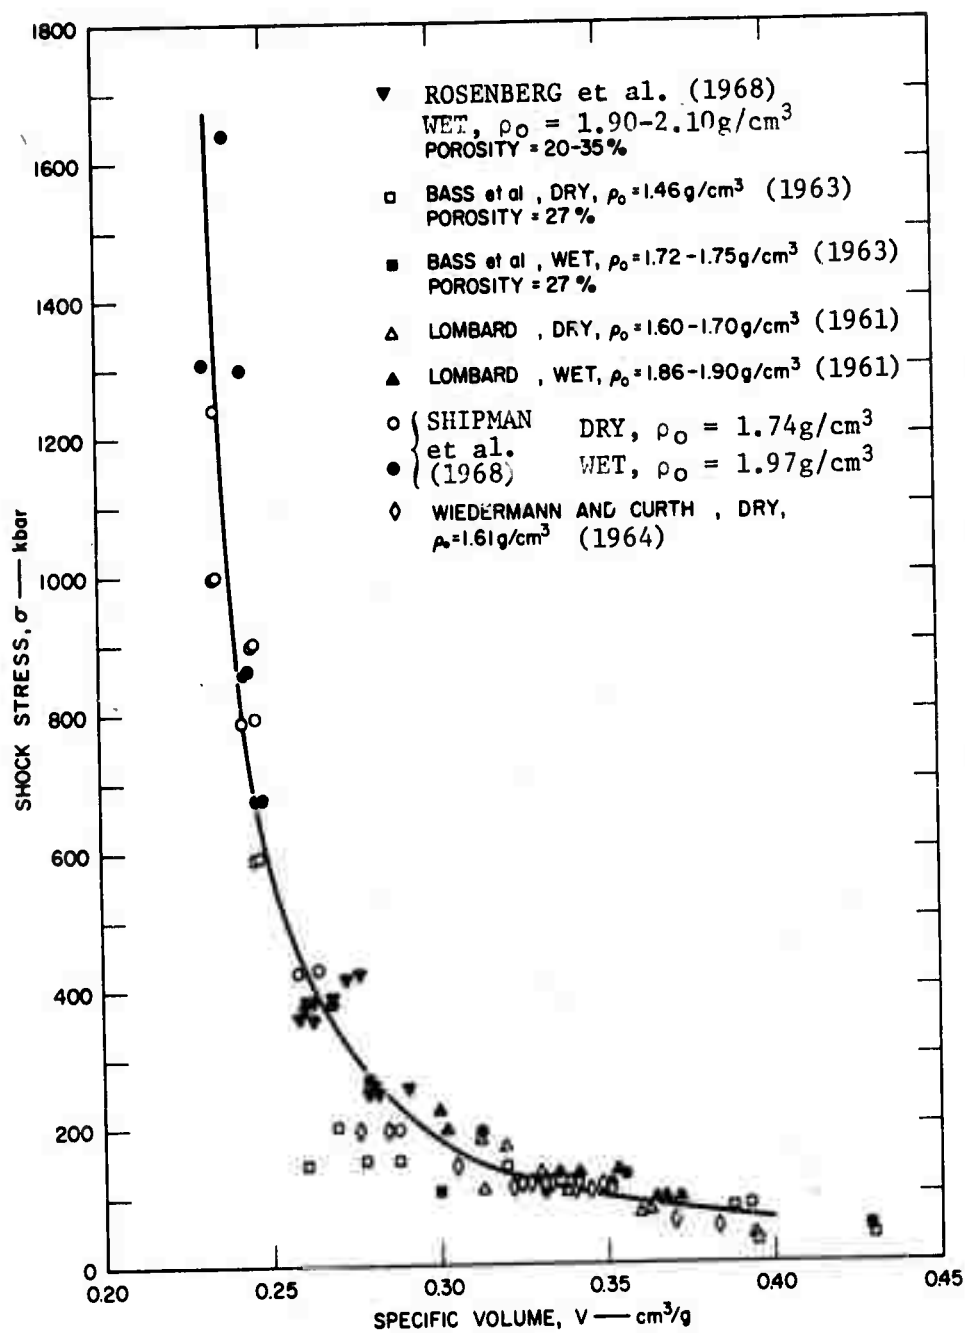
Figure 77. Hugoniot Data for Various Dry and Wet Volcanic Tuffs.
(After Rosenberg et al., 1968).

169
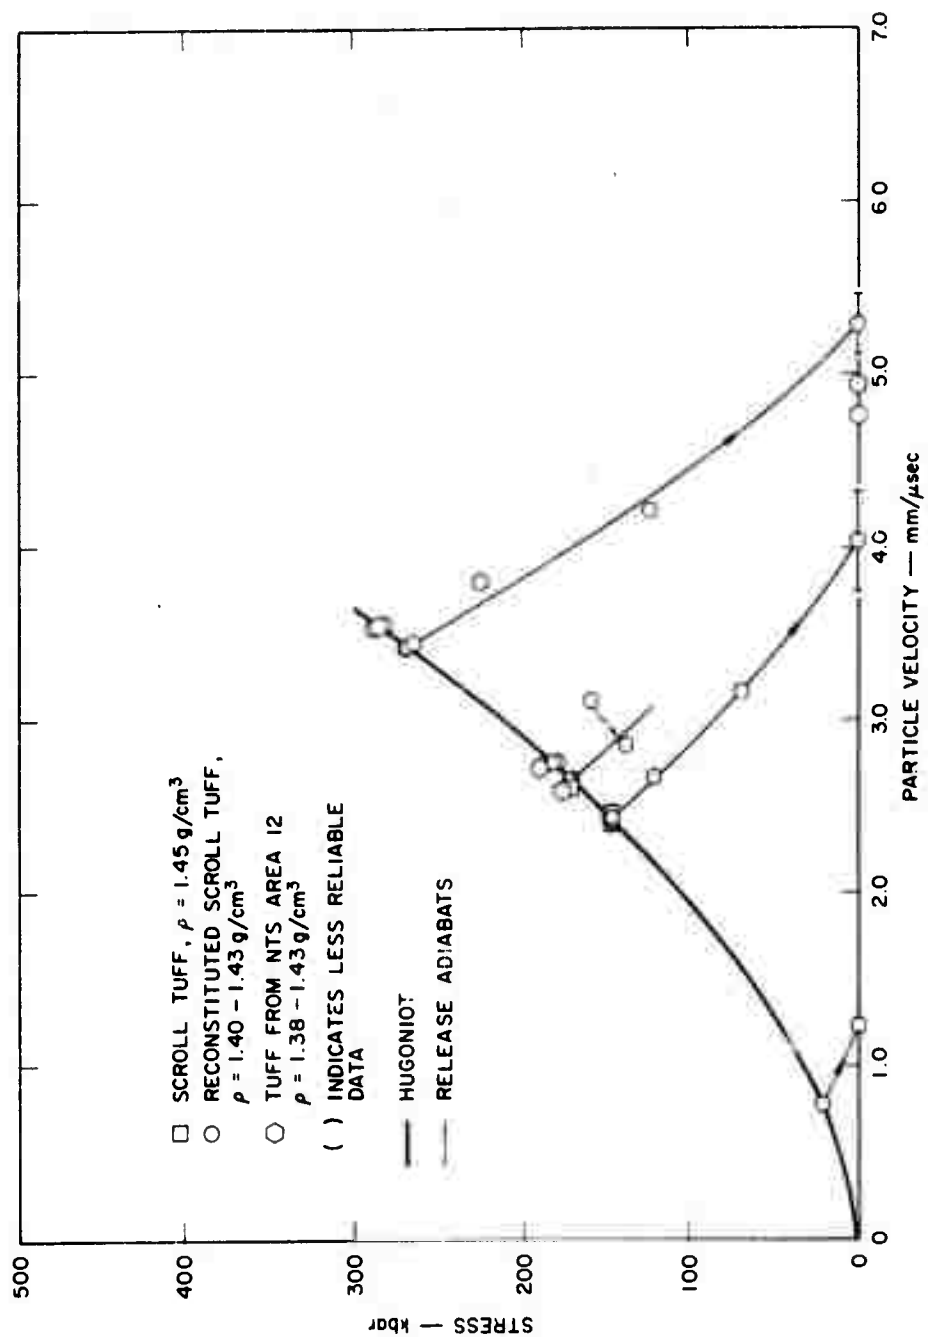
Figure 78.   Hugoniot and Release Data for Scroll Tuff in Stress-Particle Velocity Plane.   Release curves are steeper than Hugoniot. (After Petersen et al., 1968).

170
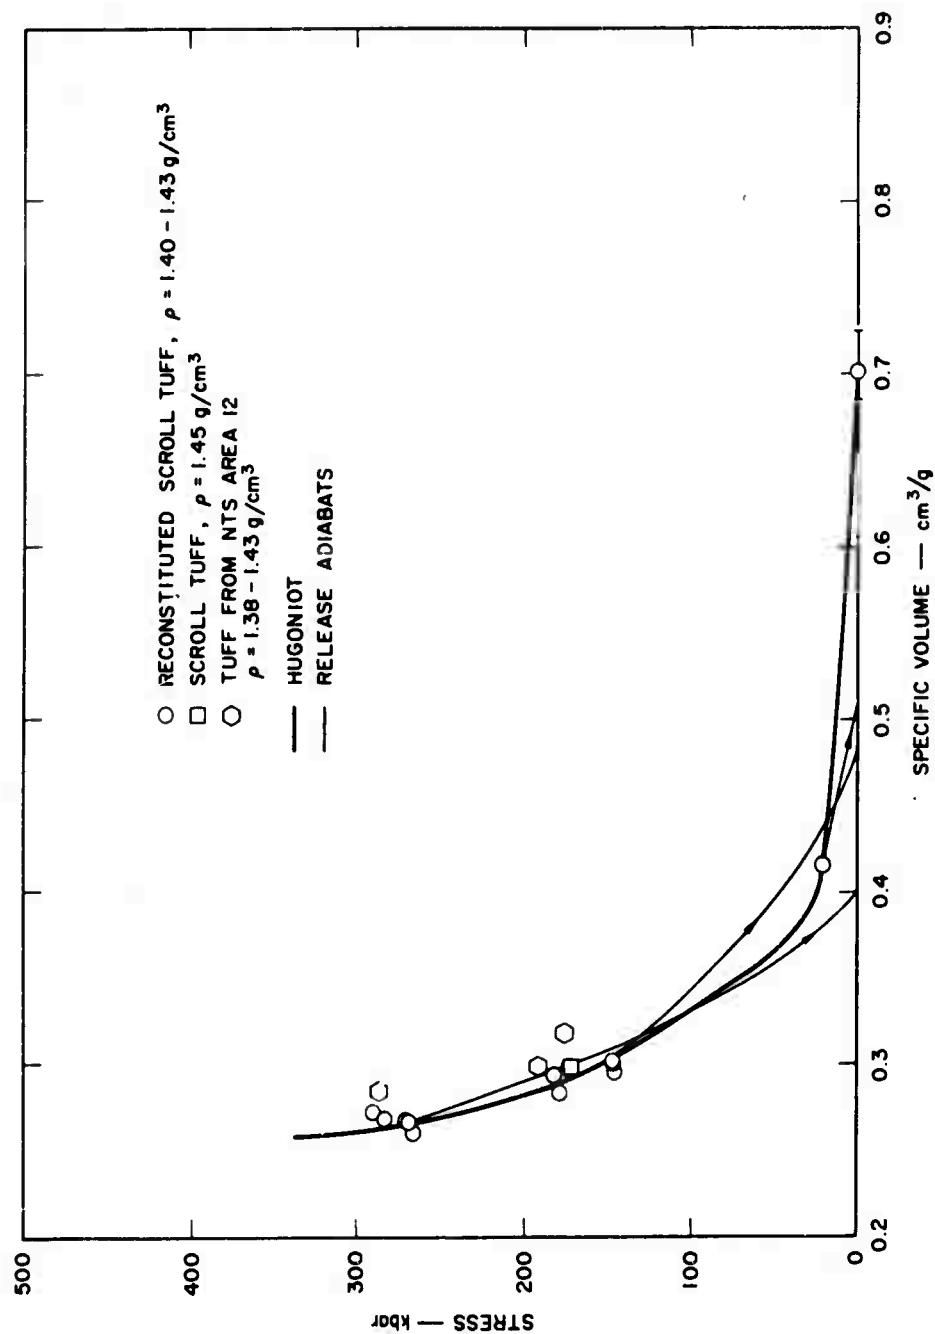
Figure 79. Hugoniot and Release Data for Scroll Tuff in Stress-Volume Plane. Release curves computed from those in Figure 78. (After Petersen et al., 1968).

171
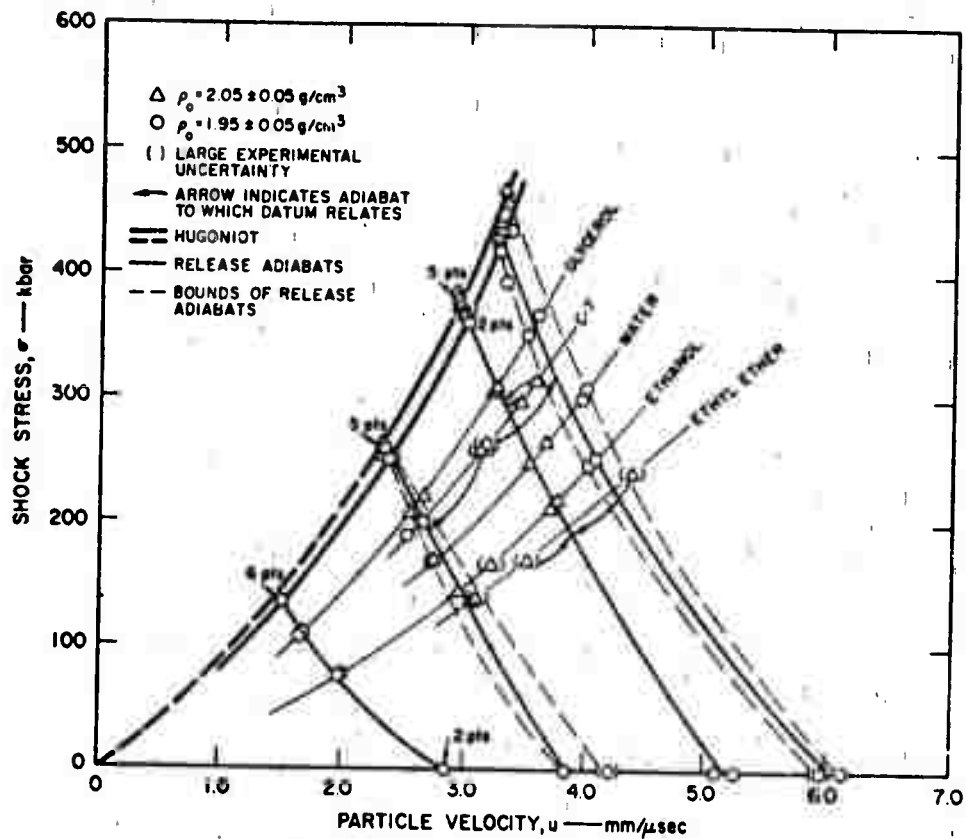
Figure 80. Shock Stress-Particle Velocity Hugoniot and Release Adiabat Data for Saturated Tuff. Two Hugoniots are drawn to emphasize two specimen density ranges. In the stress range where no measurements were made, the Hugoniot is represented by a dashed line. Hugoniots of buffer materials used in release adiabat measurements (on which the tuff release adiabat states must lie) are indicated by light lines. (After Rosenberg et al., 1968).
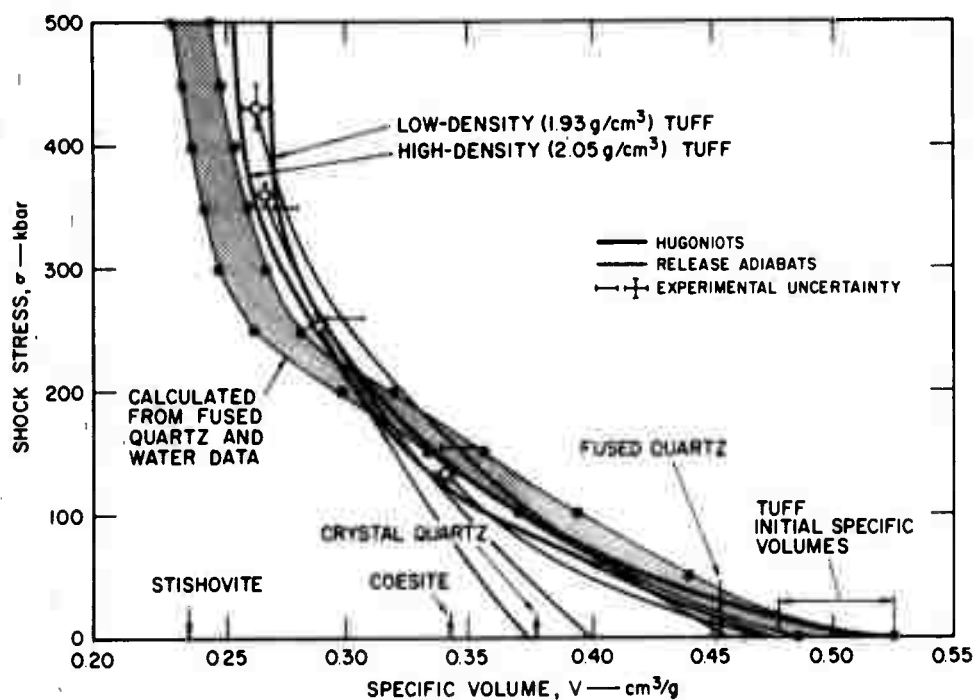
Figure 81.   Shock Stress-Specific Volume Release Adiabat Data for
Saturated Tuff.   Transformations of four tuff σ - u release
adabats of Figure 80 are shown.   Indicated bounds of the
high pressure release adiabat (reflecting experimental un-
certainty) were obtained by transforming the two extreme
fits to the σ - u data.   Standard-state specific volumes of
several phases of SiO_2 are indicated.   (After Rosenberg et
al., 1968).
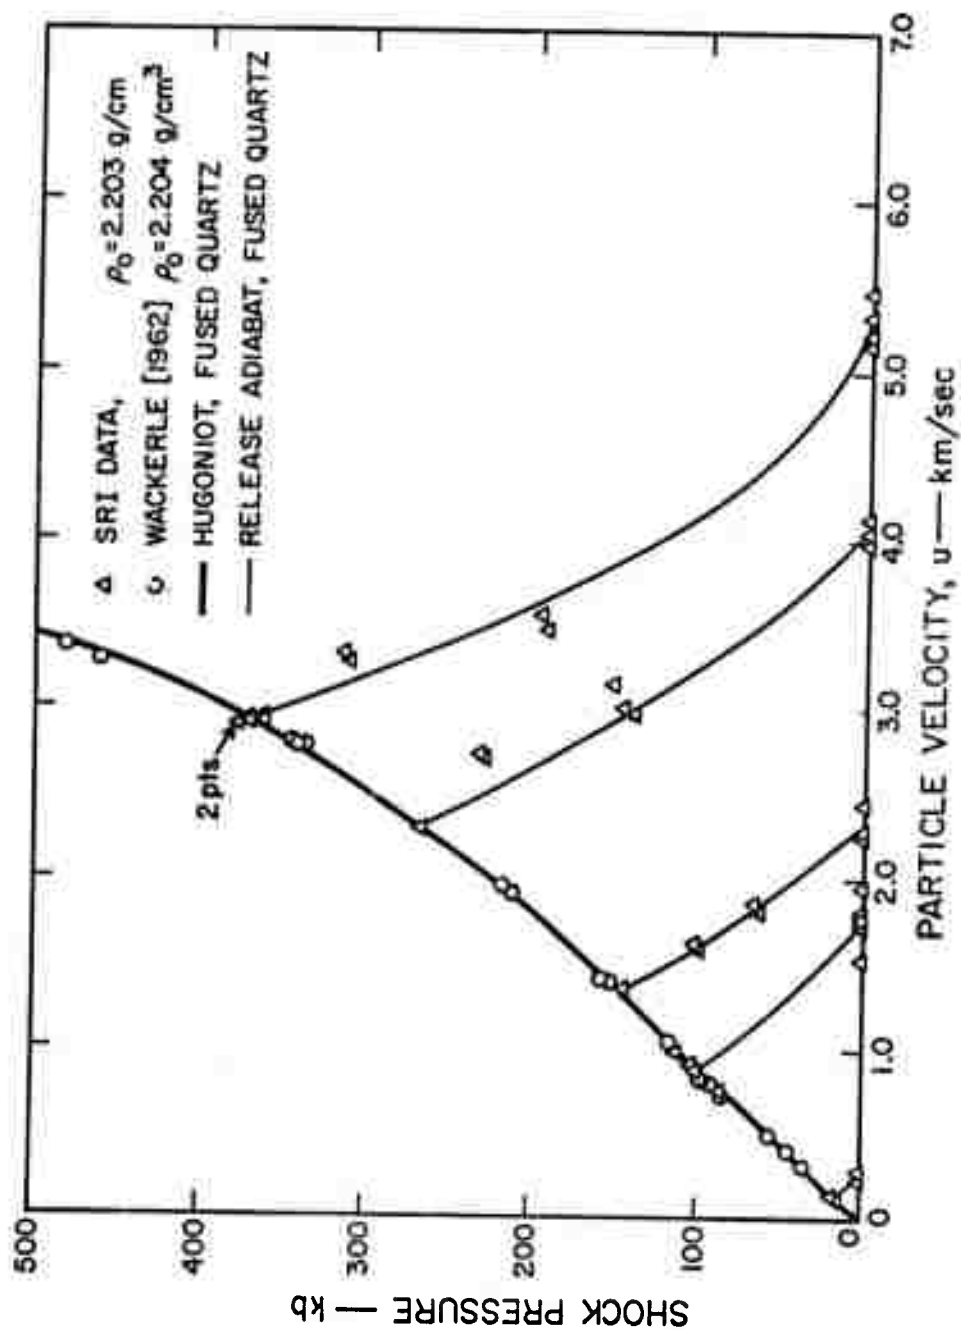
Figure 82. Shock Pressure Versus Particle Velocity, Release Adiabats for Fused Quartz. (After Rosenberg et al., 1968).

174
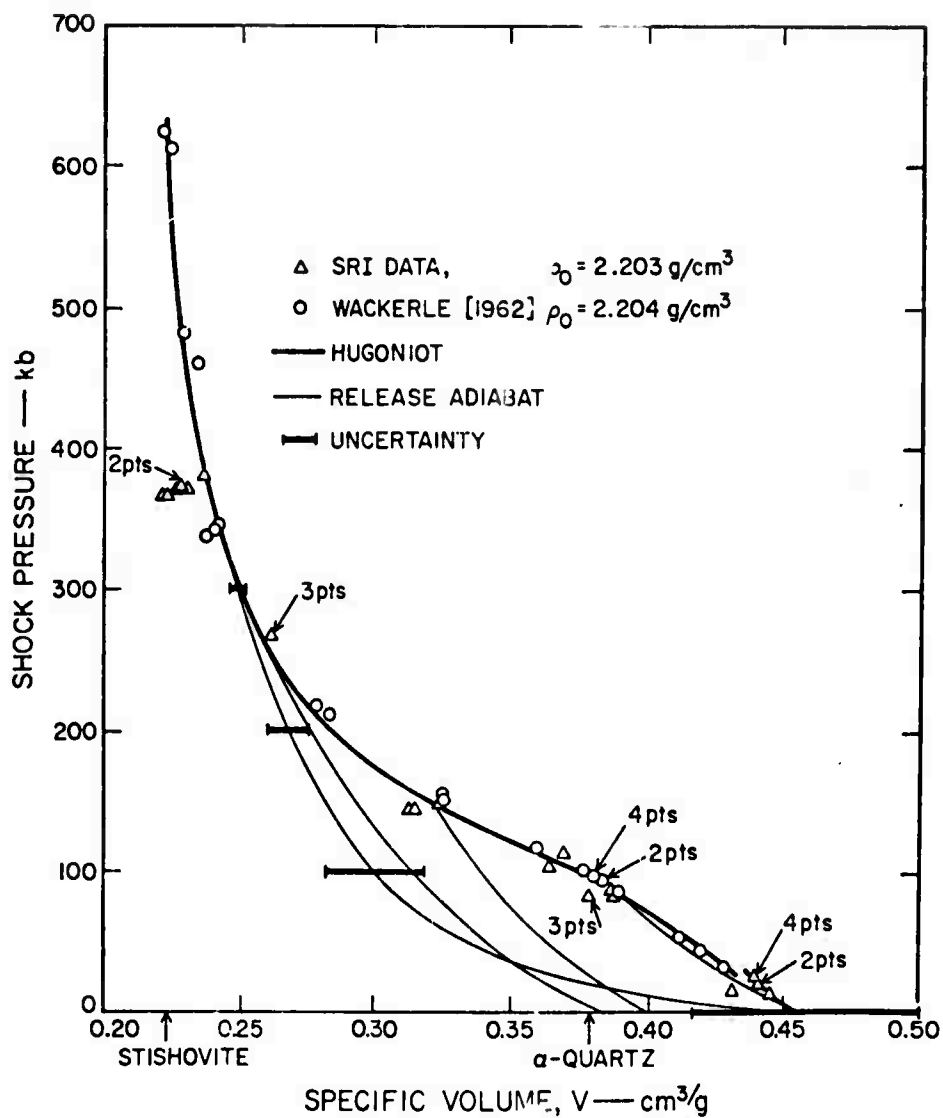
Figure 83. Shock Pressure Versus Specific Volume Release Adiabats for Fused Quartz. (After Rosenberg et al., 1968).

175
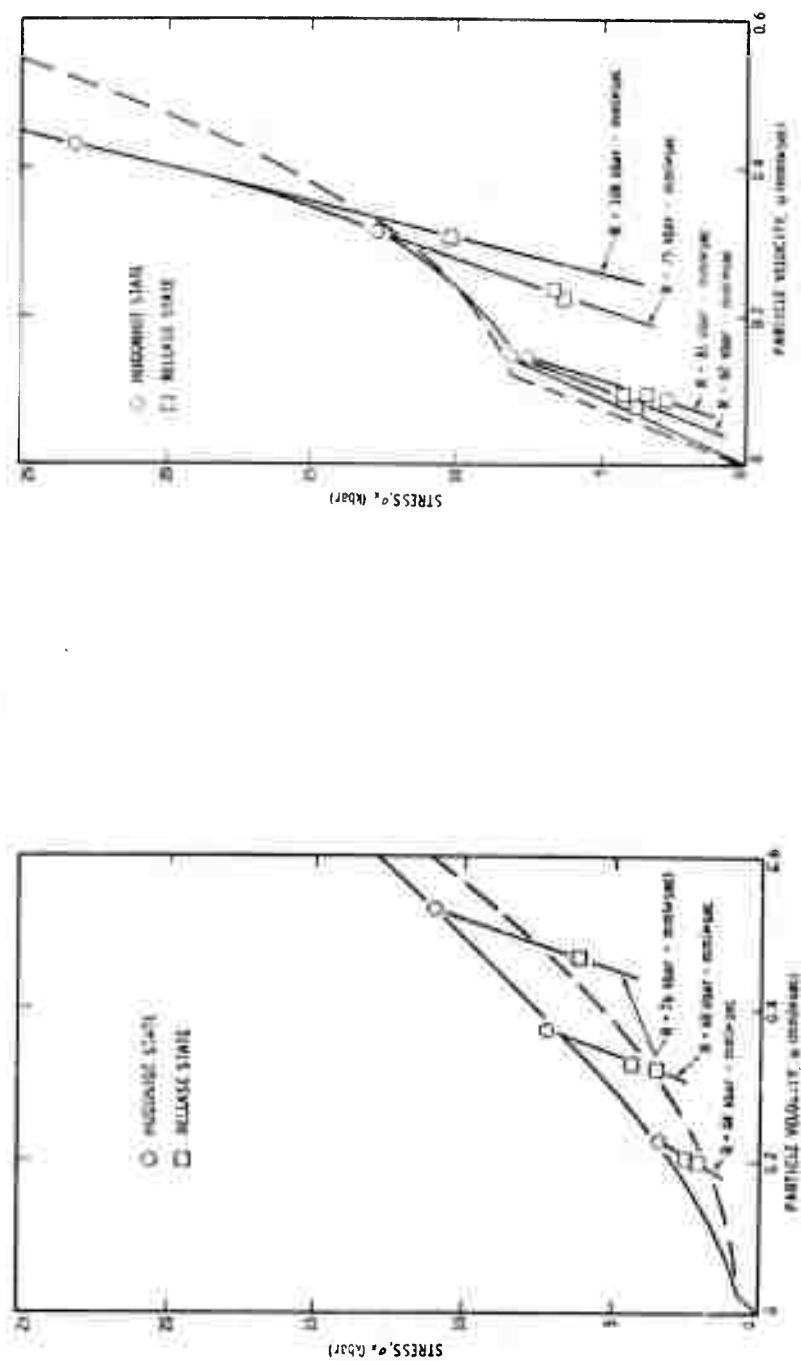
Figure 84. Hugoniot and Release Adiabats (Solid) and Calculated Hugoniot (Dashed) for the 1.34 g/cm³ Dry Tuff (Left Plot) and for the 2.00 g/cm³ Dry Tuff (Right Plot). (After Lysne, 1970).
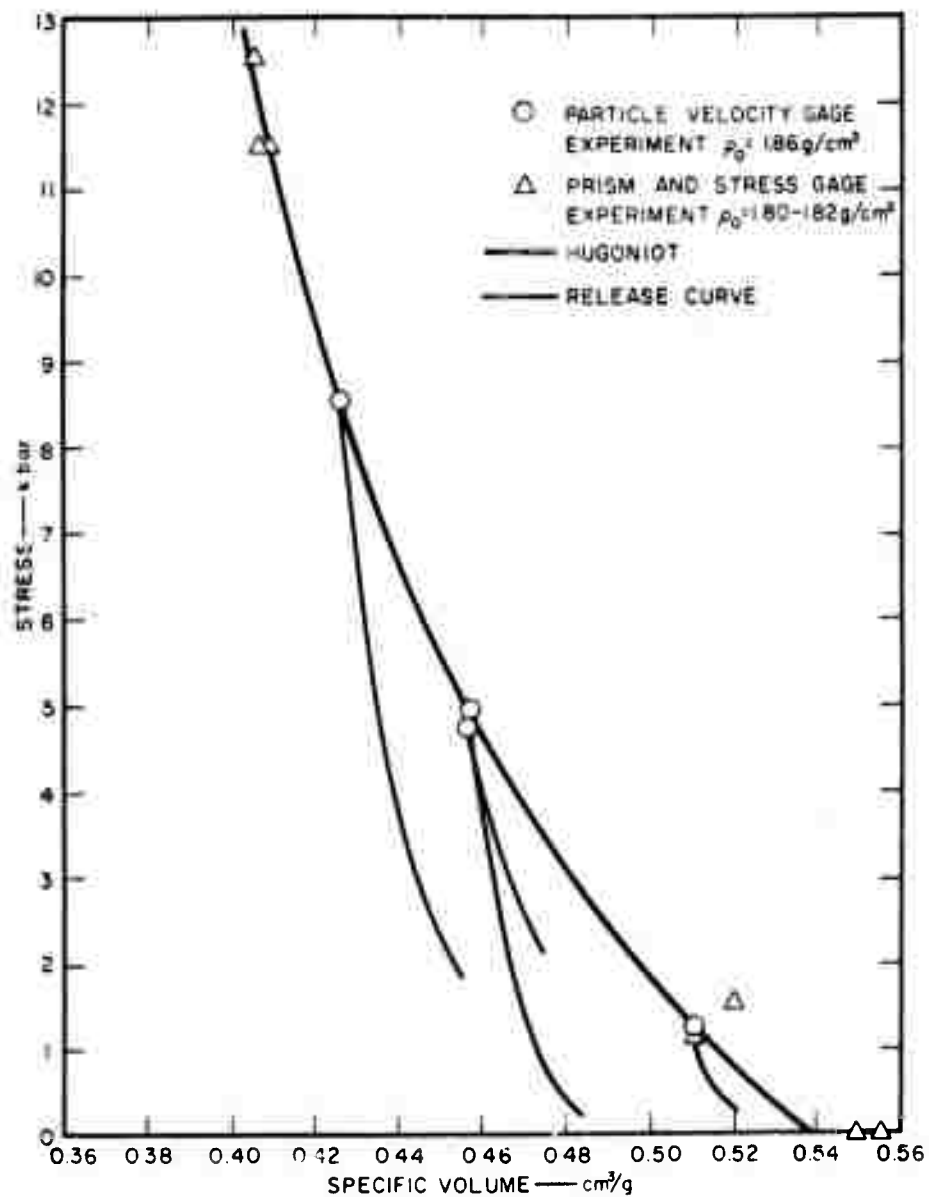
Figure 85.   Release Adiabat Curves for Alluvium Obtained Using
Embedded Foil in Magnetic Field Technique for
Measuring Particle Velocity.
(After Petersen et al., 1969).

177

III.   Effect of Water on the Equations of State

Recent studies by Wagner and Louie (1969) and Butkovich (1970) have been concerned with the effect of interstitial water in rock during stress-wave propagation from an underground explosion. In most cases, these workers have varied critical parameters in the assumed equation of state used in finite-difference calculations to match experimental observations.

Wagner and Louie have calculated particle-velocity profiles for the Hardhat and Piledriver explosions at various distances from the source using a one-dimensional elastoplastic code. By varying the rock-water content, yield strength, and failure criterion, they have attempted to fit calculated particle-velocity profiles and calculated decay of particle velocity with distance to the experimental data. In formulating an equation of state for the water-rock mixture, no irreversible phase change in the rock components was assumed. Water was assumed to have an equation of state as shown in Figure 86. Both components of the rock-water mixture are assumed to have equal pressures and temperatures during both the shock and the release process. In their calculation, the internal energy imparted to the mixture is constantly subdivided to satisfy this condition. The significance of this assumption versus the assumption that each component is shocked along its own Hugoniot and releases along its own release adiabat with no thermal flow taking place between components is discussed in a recent report by Riney et al. (1970). Assuming constant pressure and temperature in both components, Wagner and Louis obtained some surprising results in that a yield strength of only 250 bars gave calculated velocity profiles which agreed more closely with velocity gage data for the Hardhat and Piledriver explosions than the usual assumption of high yield strength for granite. In fitting experimental data describing shock wave decay from explosions in granite, similar calculations employing a dynamic yield strength of ~300 bars have been carried out by C. S. Godfrey and coworkers (private communication, 1968). They used a simple elastoplastic model and a Mohr-Coulomb-type yield criteria. More recently, McKay and Godfrey (1969) have shown that a constitutive model in which a yield strength comparable to values which are observed in the laboratory (HEL), ~50 kb, can be retained if, at later times in the flow, a strength value that is appropriate for block slippage is employed. In the block-slippage model, the dominant yielding phenomena occur during the rarefaction process and are controlled by the coefficient of friction between presumably wet rock blocks. With this constitutive model, McKay and Godfrey are able to closely fit Hardhat and Piledriver stress wave data (Figures 87-88).

Butkovich has employed a one-dimensional (SOC) finite-difference code to calculate the radial-stress profiles and free-surface velocity profiles from explosions in water-saturated tuff and granite. The latter calculation was, in some respects, similar to the Wagner and Louie calculations in one dimension.
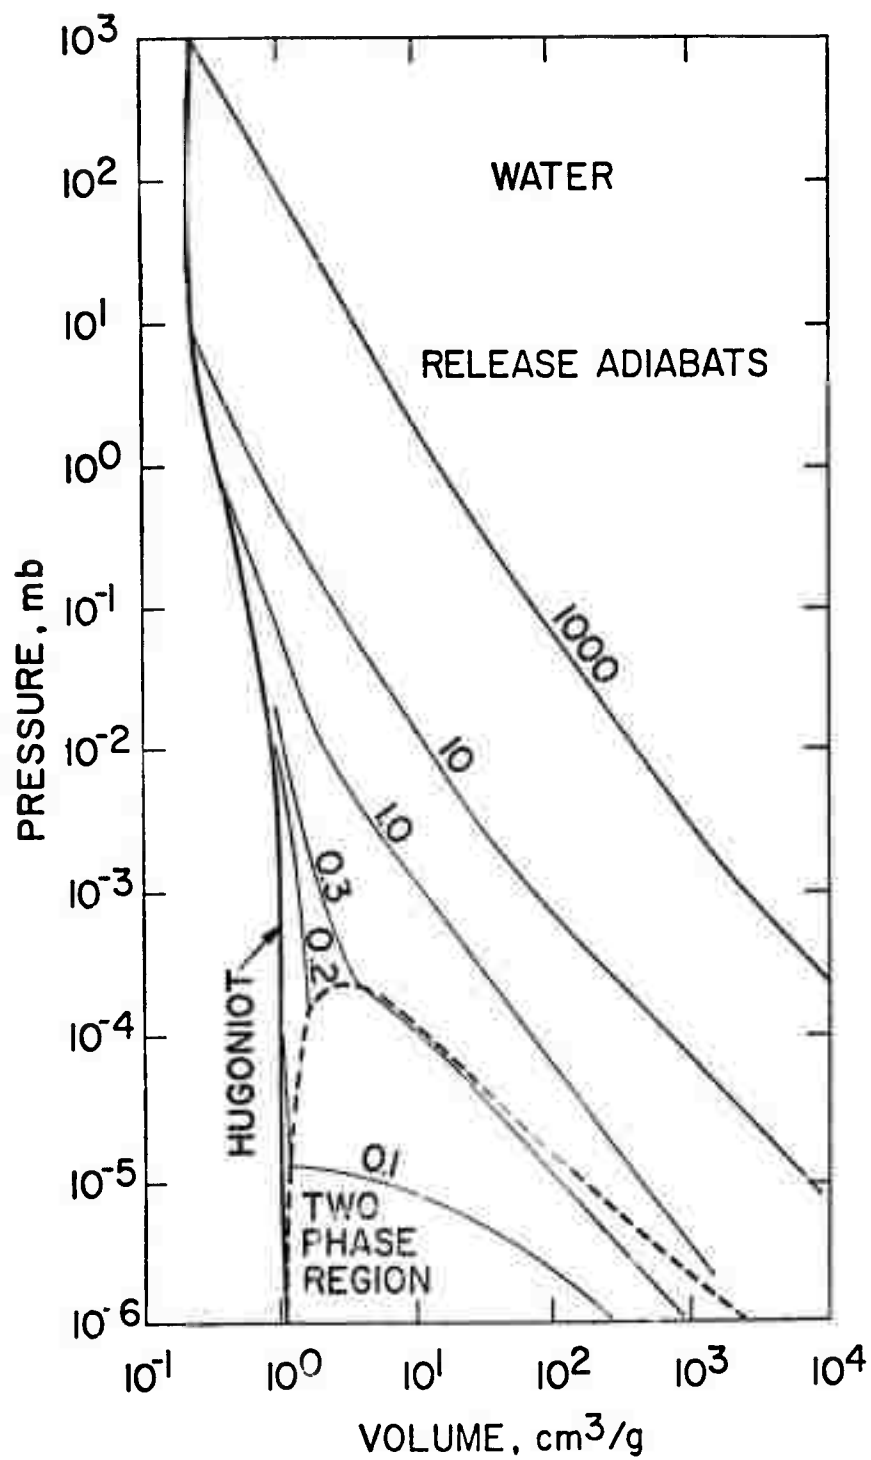
Figure 86. Hugoniot for Water and Calculated Release
Adiabats. (After Bjork et al., 1969).

179
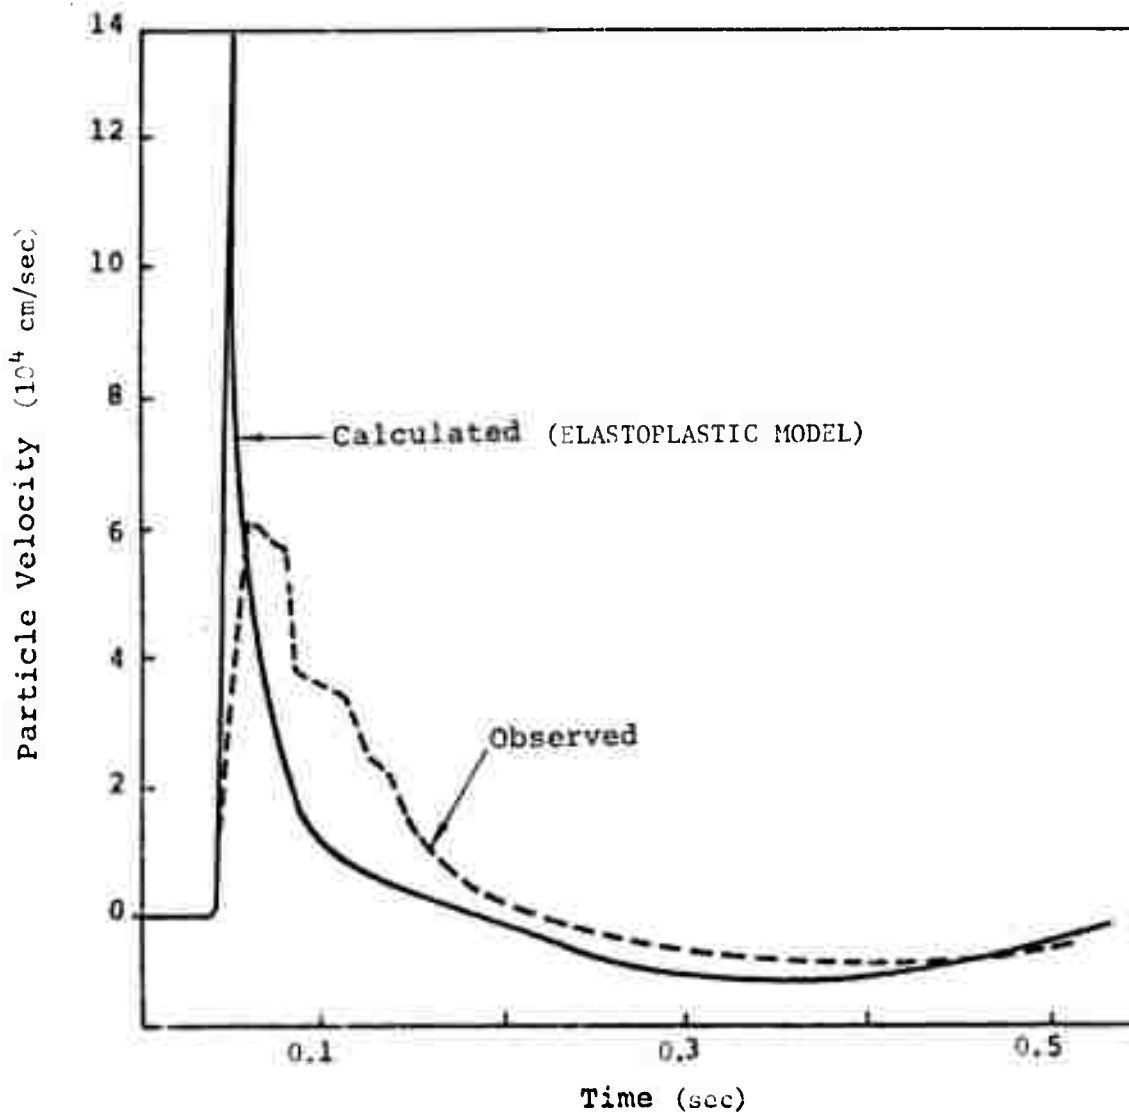
Figure 87. Particle Velocity Versus Time Profiles, 1600 ft from Shot
Point for Piledriver Explosion. Calculated profile obtained
using only elastoplastic equation of state and no block
slippage. (After McKay and Godfrey, 1969).
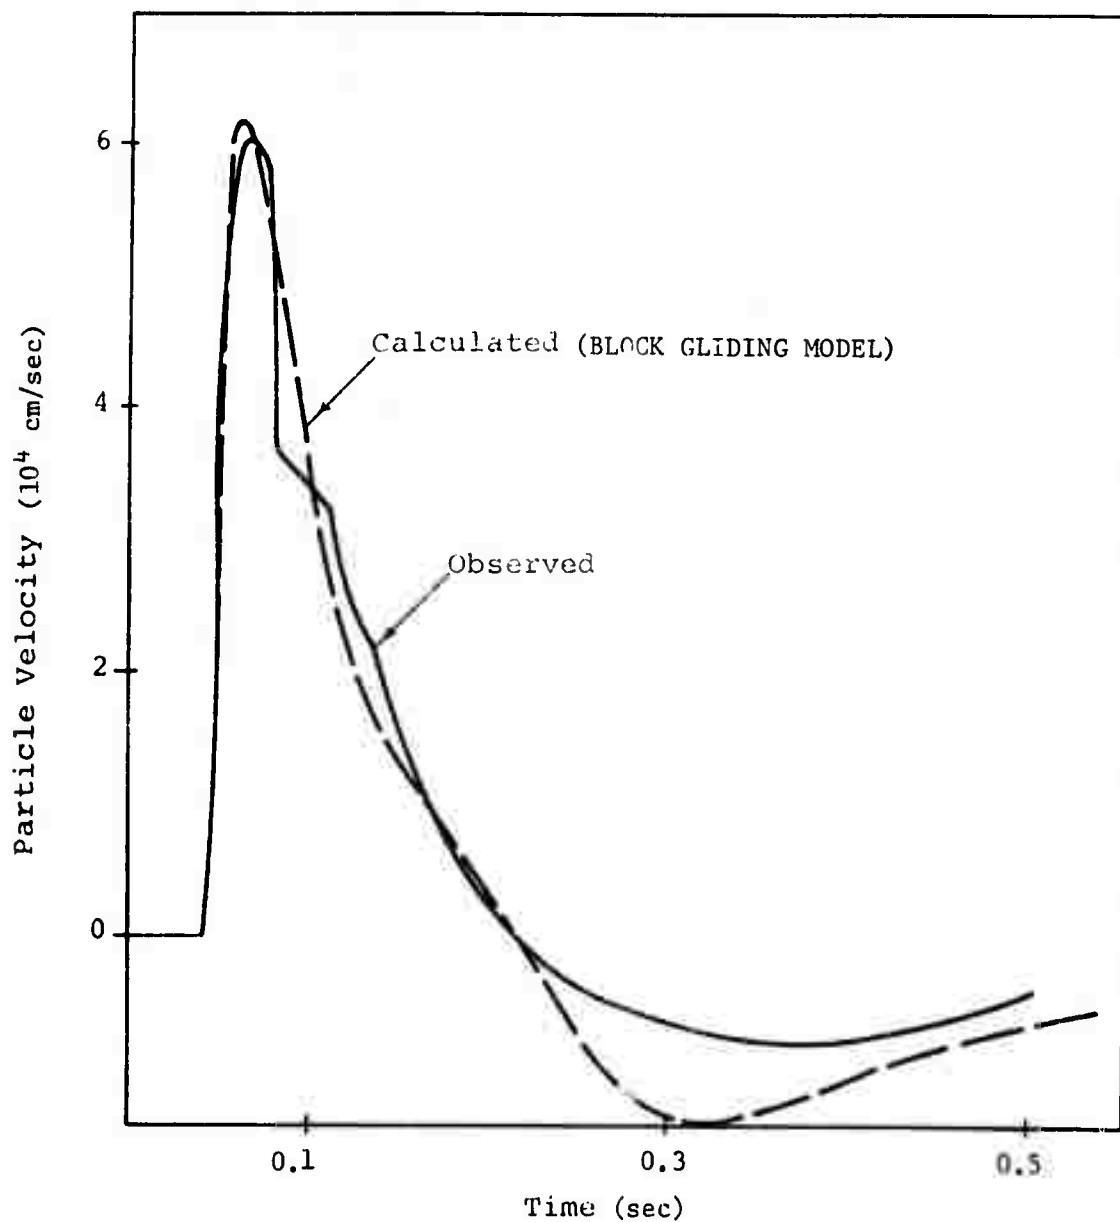
Figure 88. Particle Velocity Versus Time Profiles, 1600 ft from Shot Point for Piledriver Explosion. Calculated profile employs block gliding model. (After McKay and Godfrey, 1969).

181

Although the same equation of state for water was used, it was combined with that of dry rock in a different way. Butkovich assumes that both materials are shocked to the same final pressure (but not temperature) state and this state corresponds to a mass average of the principal Hugoniot for each material. Adiabatic release occurs separately for each component and an average release adiabat is calculated from simple mixture theory (Figure 89). The marked effect of including water in these calculations on the release portion of the stress-wave profile is shown in Figures 90 and 91. In several of Butkovich's calculations, the irreversible phase change of the silicate component to a dense stishovite-like phase was also assumed (Figure 92). The results from this refinement have tended to increase the agreement between the field data for the Benham explosion and the calculations as shown in Figures 93 and 94.

IV. New Equation of State Data in the High Pressure Regime

A considerable body of high-pressure Hugoniot, and in a few cases release adiabat data, pertinent to the description of stress-wave propagation in the intense stress, or high pressure, regime has recently been reported. At the previous VELA Uniform Point Source Meeting, McQueen (1968) reported Hugoniot data for a wide class of rocks and minerals (mostly silicates). Some of these data are reported in Clark (1966), as well as in McQueen et al. (1967). Virtually all of the silicates and several of the oxides display Hugoniot curves such as shown in generalized form in Figure 95. Much of these data have been analyzed and interpreted by a series of workers, including McQueen et al. (1967), Anderson and Kanamori (1968), Wang (1967), Ahrens et al. (1969a), Ahrens et al. (1970), and more recently by Davies and Anderson (1971). As indicated in Figure 95, the Hugoniot data suggest the existence of at least three regimes: a low-pressure regime in which the Hugoniot represents the shock equation of state of the initial low-pressure phase; a mixed-phase regime, in which the Hugoniot states represent a mixture of the low-pressure and high-pressure, shock-induced phase; and finally, a high-pressure regime in which the Hugoniot equation of state represents the properties of a wholly transformed, denser, shock-induced phase. Of some 25 silicates, carbonates, and oxides that have been studied to pressures of approximately 1 mb, only MgO, Al_2O_3, and MnO_2 appear to remain in their low-pressure phase over the range of pressures that has been investigated. Virtually all the rocks investigated, including ones with high porosity, undergo at least one shock-induced phase change over the pressure range that has been explored. It should be further pointed out that aside from porosity effects, the density increase accompanying these phase changes varies from approximately 10 percent for minerals such as enstatite ($MgSiO_3$) and almandine (($Fe,Mg)_3Al_2Si_3O_{12}$), to ~60 percent for quartz. Hence, these phase changes, which in many of the silicates represent a change in the silicon-oxygen coordination from four to six, account for most of the compression which takes place in the rock upon shock compression up to 1 mb. Analysts of the high shock pressure data (mentioned above) recently have largely concentrated their efforts on estimating the zero-pressure density, complete equation of state, and possible crystallographic structure of the high-pressure phases of the silicates.
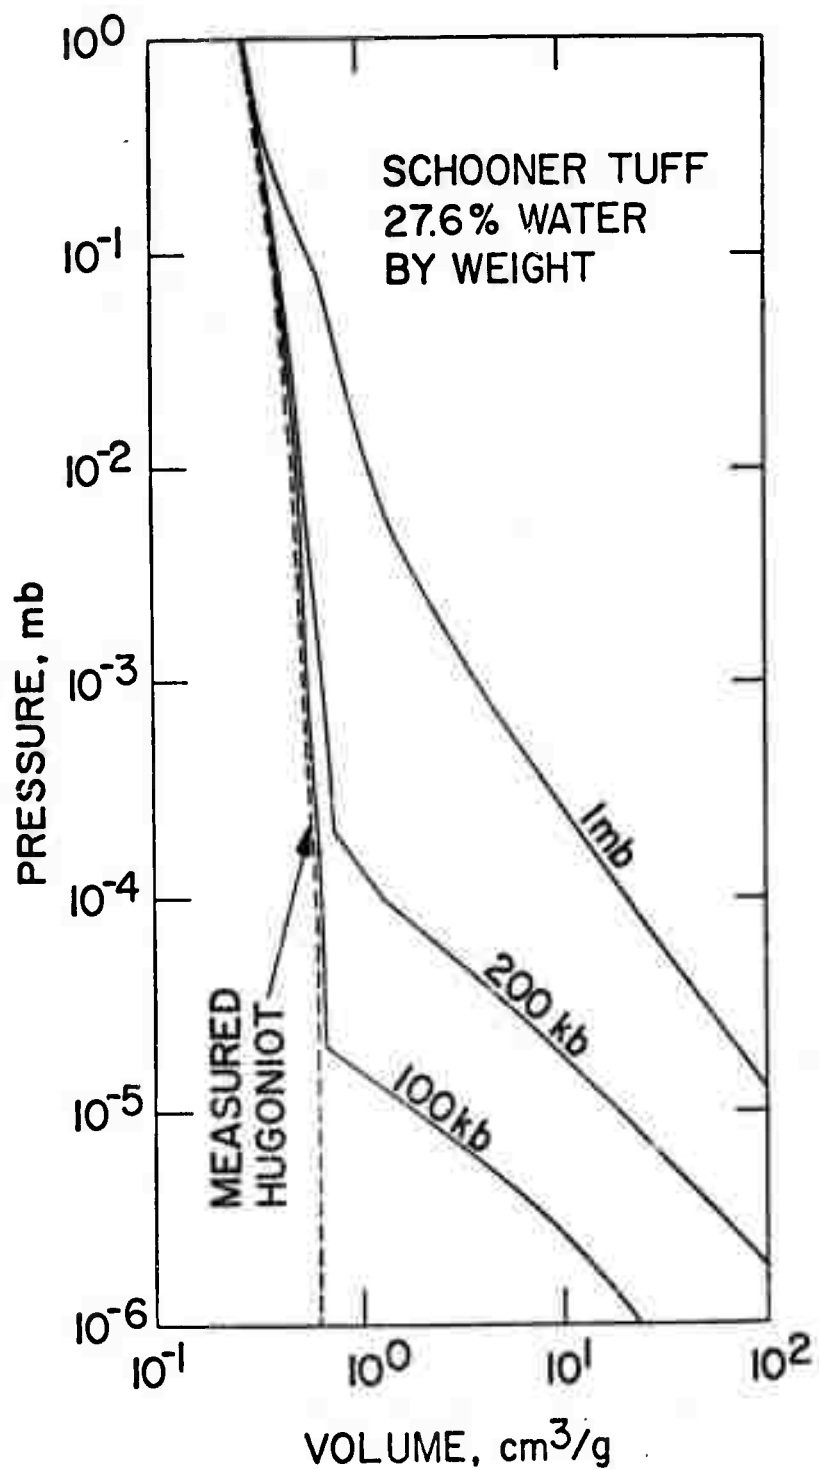
Figure 89. Hugoniot and Calculated Release Adiabats for Water-Saturated Schooner Tuff. (After Butkovich, 1970).

183
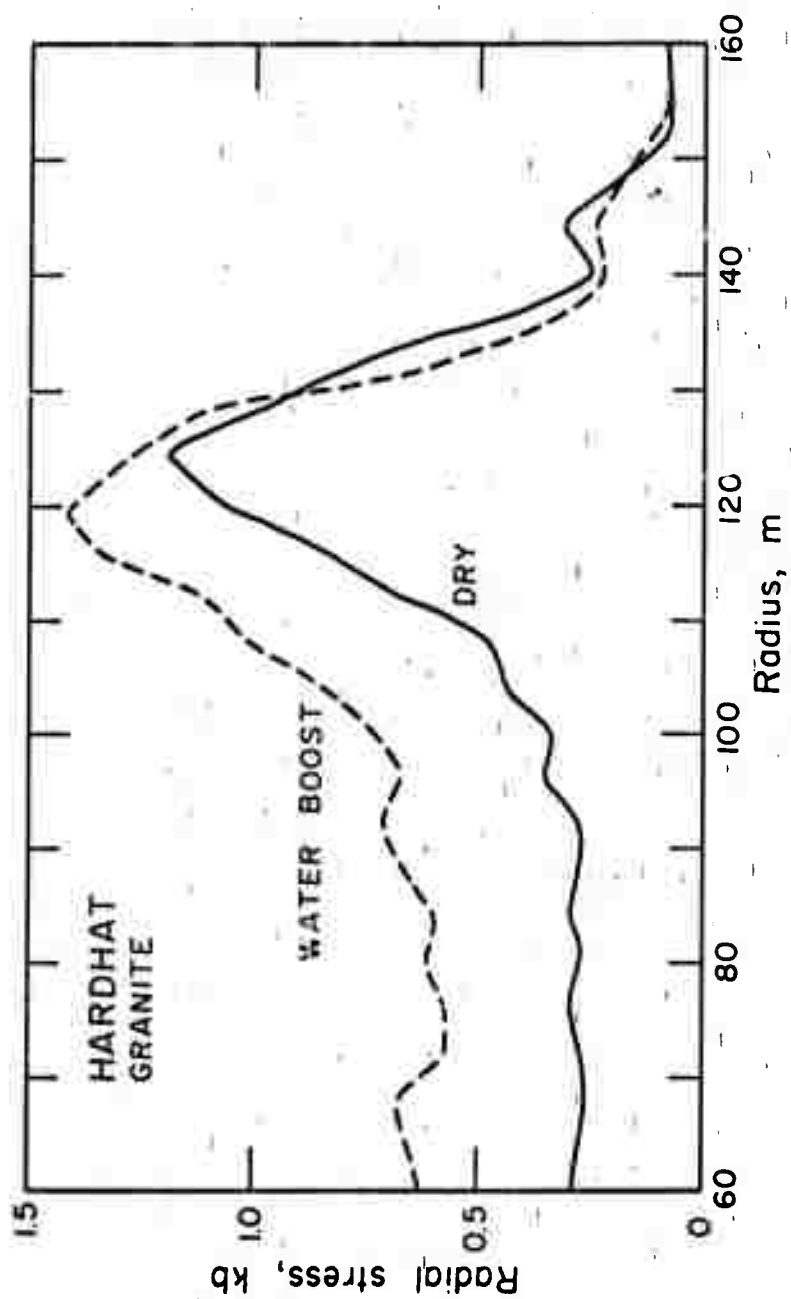
Figure 90. Radial-Stress Profile at 25 msec after Hardhat Explosion. Effect of water on the latter portions of the calculated radial stress profiles is observed. (After Butkovich, 1970).

184
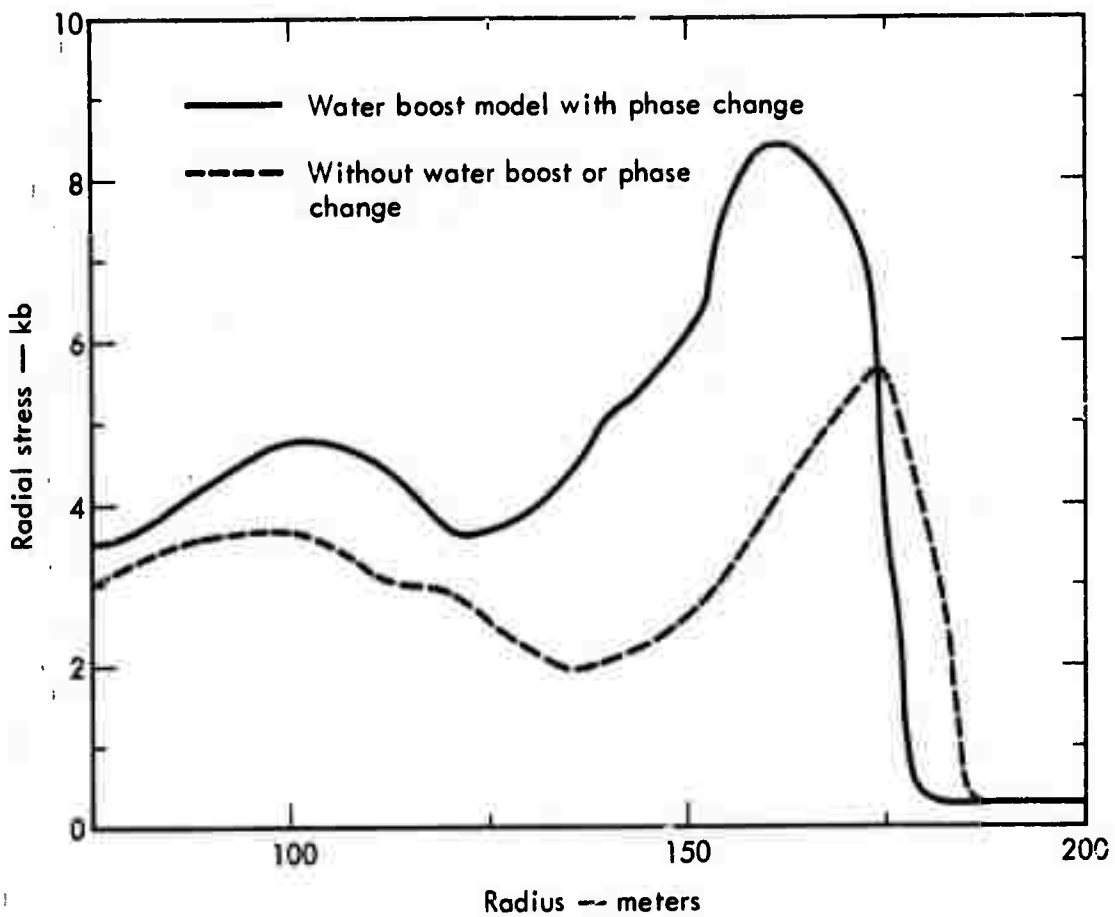
Figure 91. Radial-Stress Profiles at 50 msec in Tuff after Benham Explosion. Effect of water boost and phase change on calculated stress wave is shown. (After Butkovich, 1970).
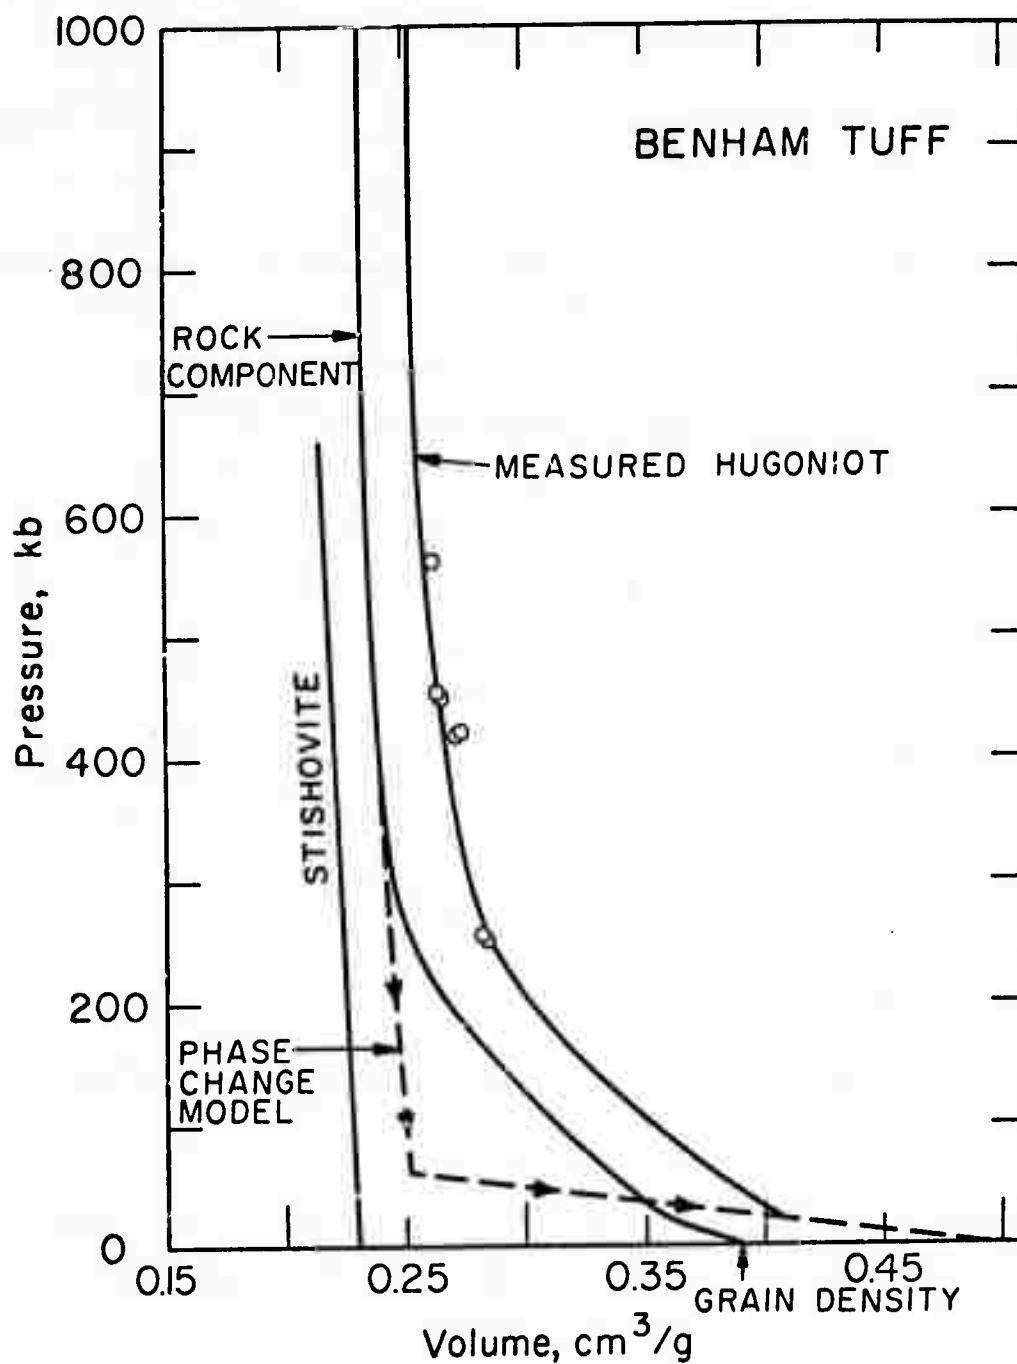
Figure 92. Hugoniot for Benham Tuff Containing 13 Percent Water. Release paths for silicate component of rock used in calculation to obtain profile in Figure 33 is indicated. (After Butkovich, 1970).

186
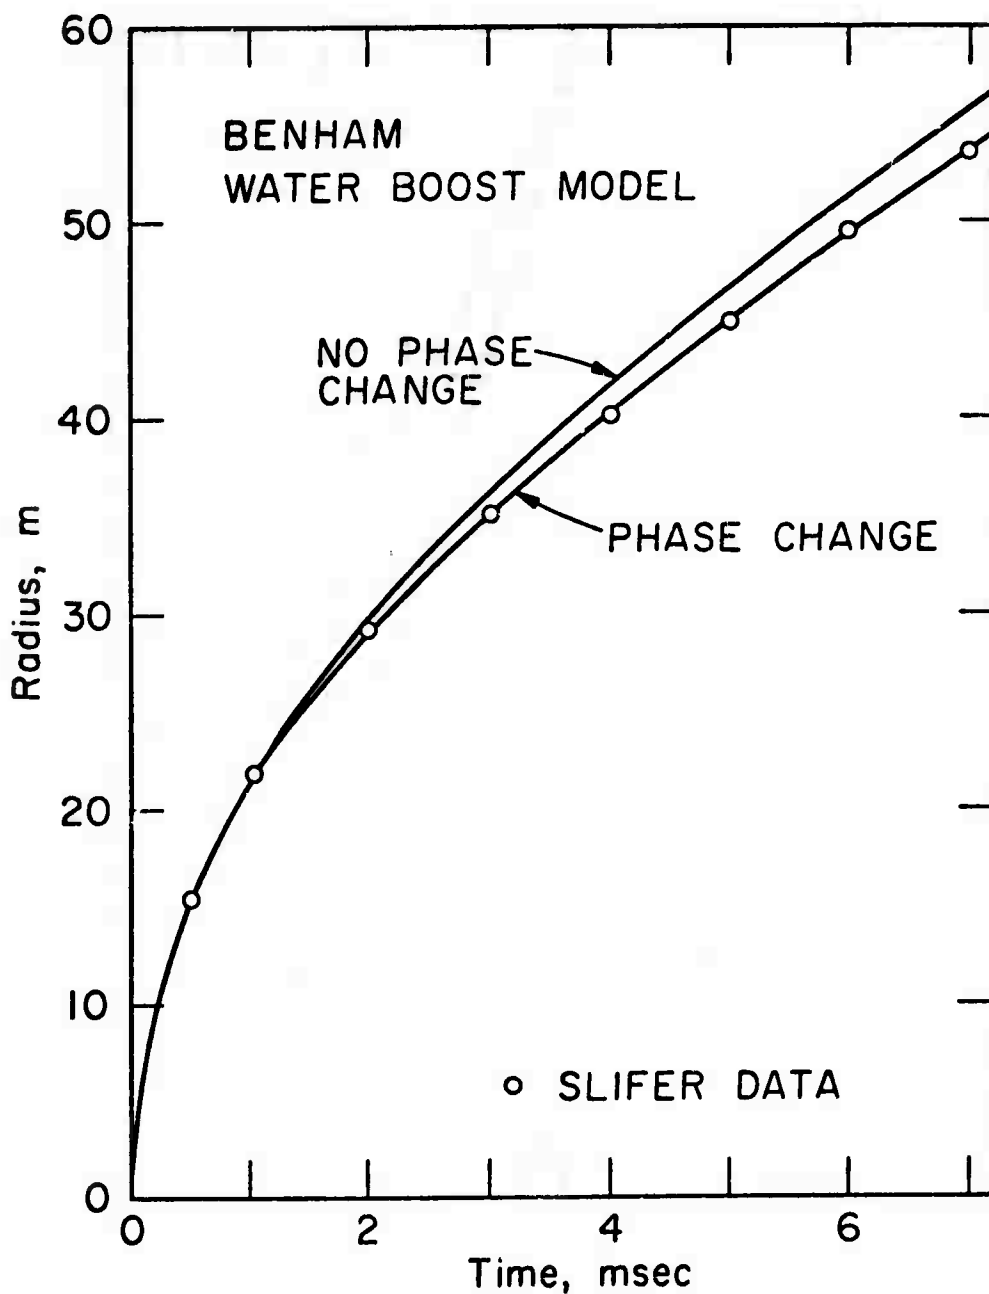
Figure 93. Shock-Wave Time-of-Arrivals, Benham Event. Calculated curves show the effect of including a phase change in the stress wave calculation. Both curves take water into account. (After Butkovich, 1970).

187
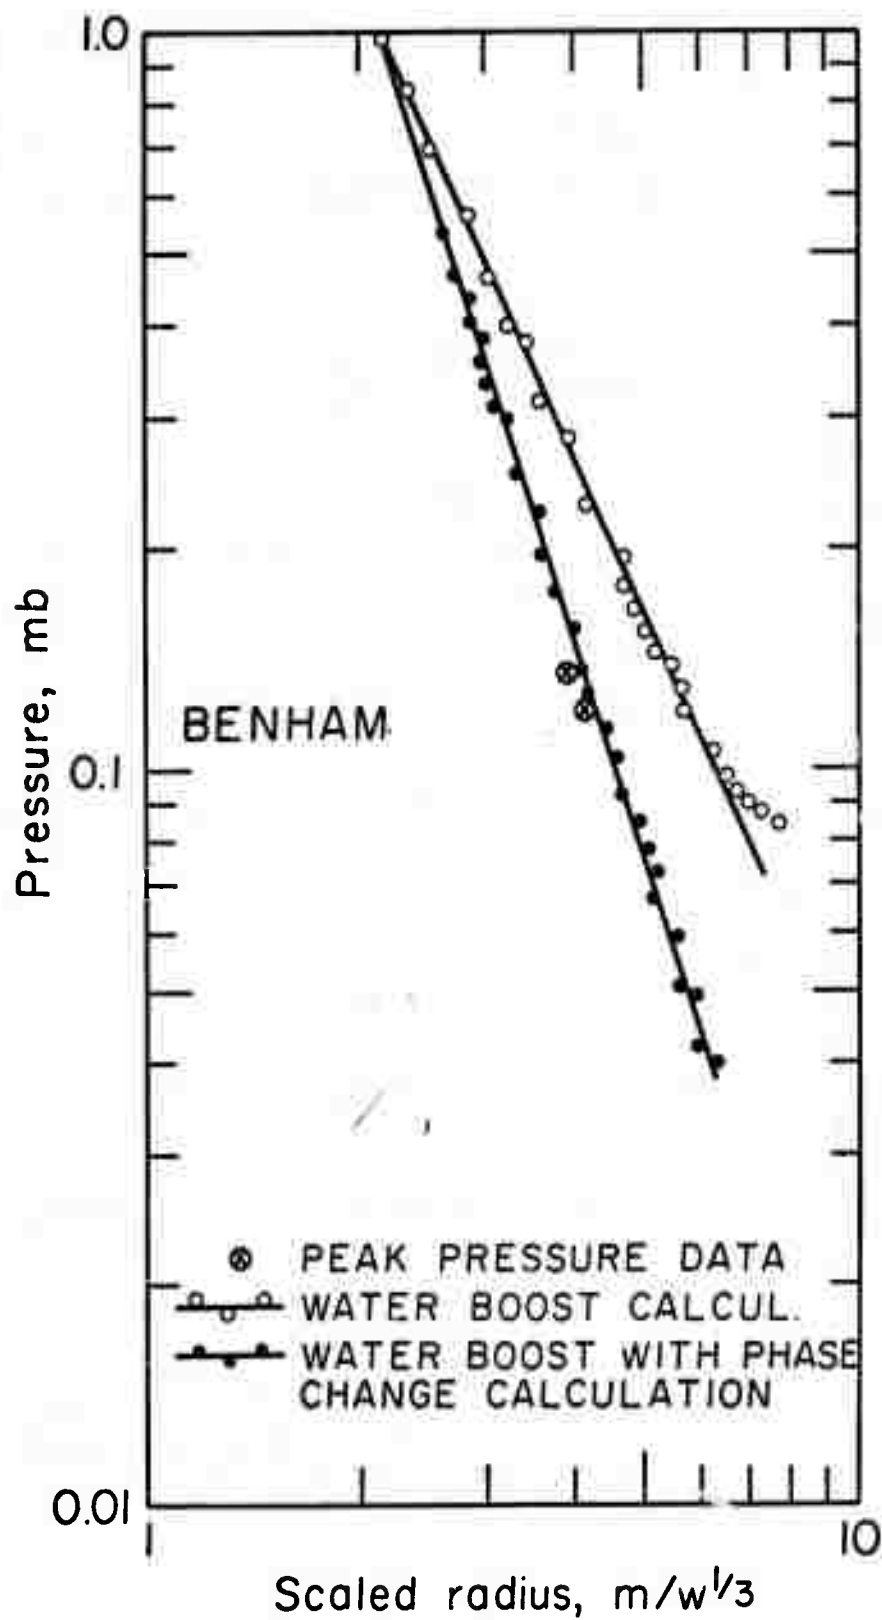
Figure 94. Calculations of Peak Pressure Versus Scaled Radius for Benham Event. Effect of including phase change in silicate component equation of state is demonstrated. (After Butkovich, 1970).
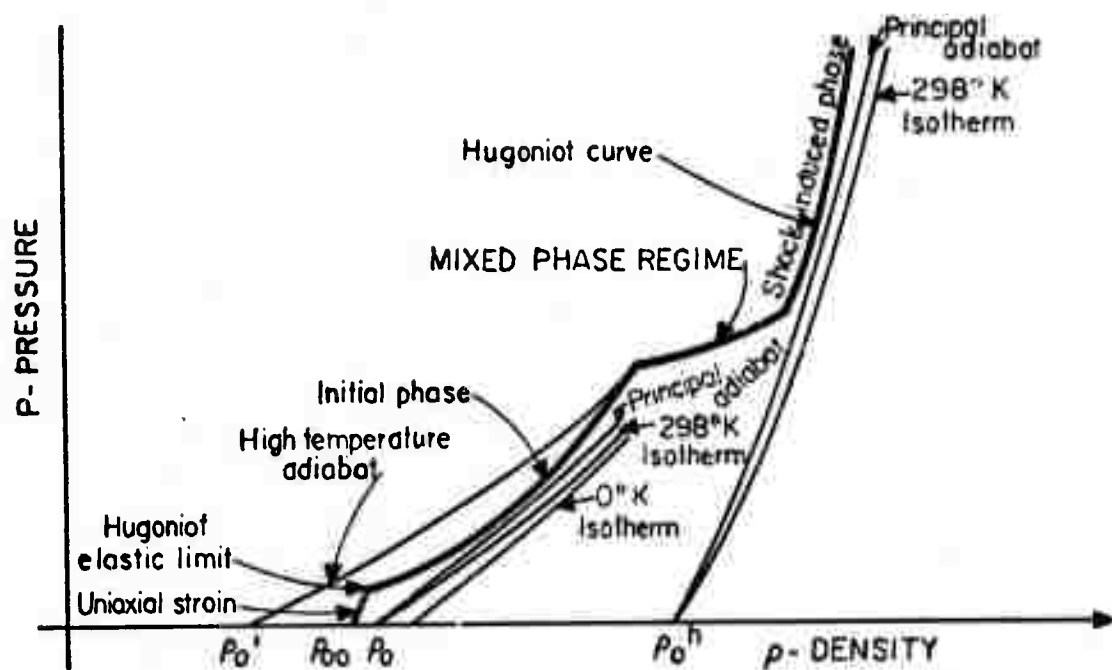
Figure 95. Generalized Hugoniot and Other Thermodynamic Curves for Silicate Mineral or Silicate-Bearing Rock.

With the exception of stishovite (SiO_2) and majorite $(Mg,Fe)SiO_3$ (garnet structure), where the natural shock phases have been recovered, the exact crystallographic nature of these phases is yet unknown.

Recently, new Hugoniot data for a series of rocks have been reported (Table 2).

Table 2. New High Pressure Hugoniot Data[a]

Material	Pressure Range Studied (kb)
Vacaville basalt	350-2000 (Figure 96)
Kaibab limestone	300-1100
Coconino sandstone	100-1400
Mono Lake pumice	40-260
Climax granite	600-2000 (Figure 97)
Dry Rainier Mesa tuff	200-1300
Wet Rainier Mesa tuff	80-1400
Gneiss	80-2100
Dolomite	100-1700 (Figure 98)
Limestone	150-1200
Alluvium	50-300

[a]Shipman et al. (1968) and Jones et al. (1968).

These results have markedly extended our knowledge of the high-pressure regime, particularly for porous rocks.

Some new release adiabat data for granite, plagioclase, and fused quartz, centered at the Hugoniot state largely in the mixed-phase region, have been reported. Release data for the Hardhat and Raymond granites were measured by Keough and Wilkinson (1967) and by Petersen et al. (1968) (Figures 99 and 100) in the range of 100 to 300 kb. Release data for plagioclase in the range 200 to 400 kb and for fused quartz in the range 150 to 300 kb have been reported by Ahrens et al. (1969b) and by Rosenberg et al. (1968) (Figures 101-102 and 82-83). Generally these data indicate that upon adiabatic release the silicates show irreversible compaction due to phase change to a denser material for shocks up to 200 or 300 kb, depending upon the mineral.
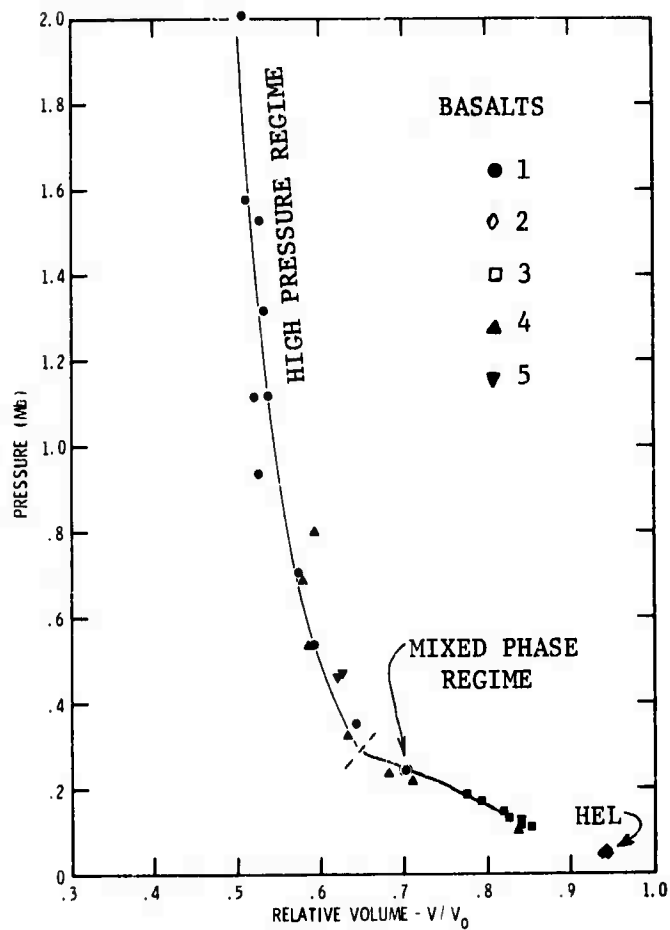
Figure 96. Hugoniot Data for Various Basalts. Demarcation of mixed phase and high pressure regimes are shown. (1) Vacaville basalt, Shipman, et al., 1969; (2) elastic shock; (3) deformation shock, Vacaville basalt, Ahrens and Gregson, 1964; (4) NIS basalt, Van Thiel et al., 1969; (5) NTS basalt, Bass et al., 1963. (After Jones et al., 1968).
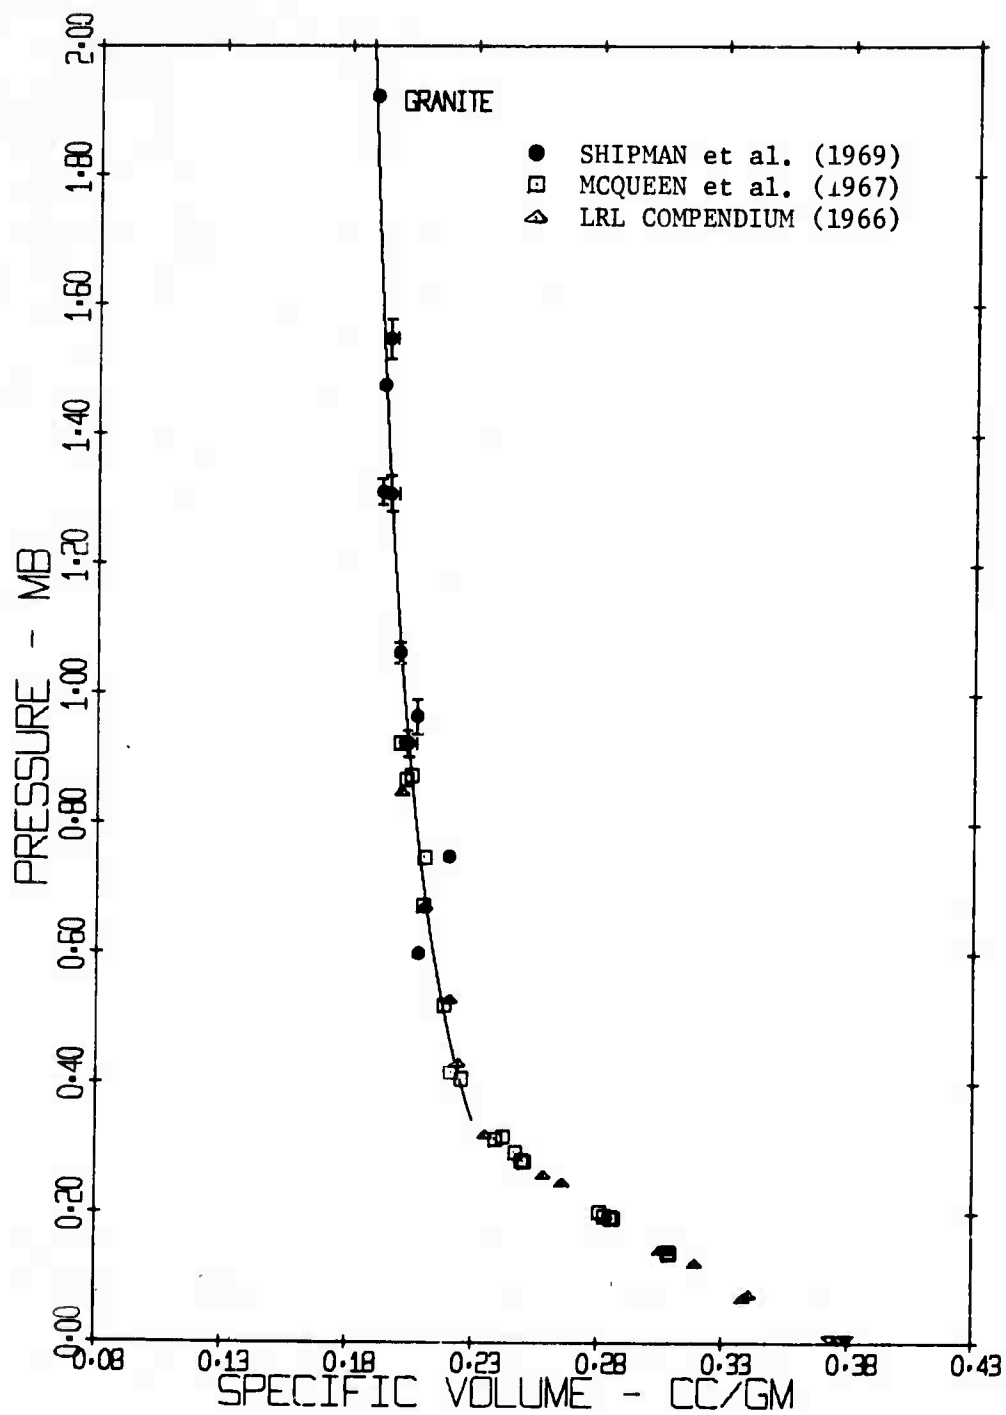
Figure 97. High-Pressure Hugoniot Data for Various Granites. (After Shipman et al., 1969).

192
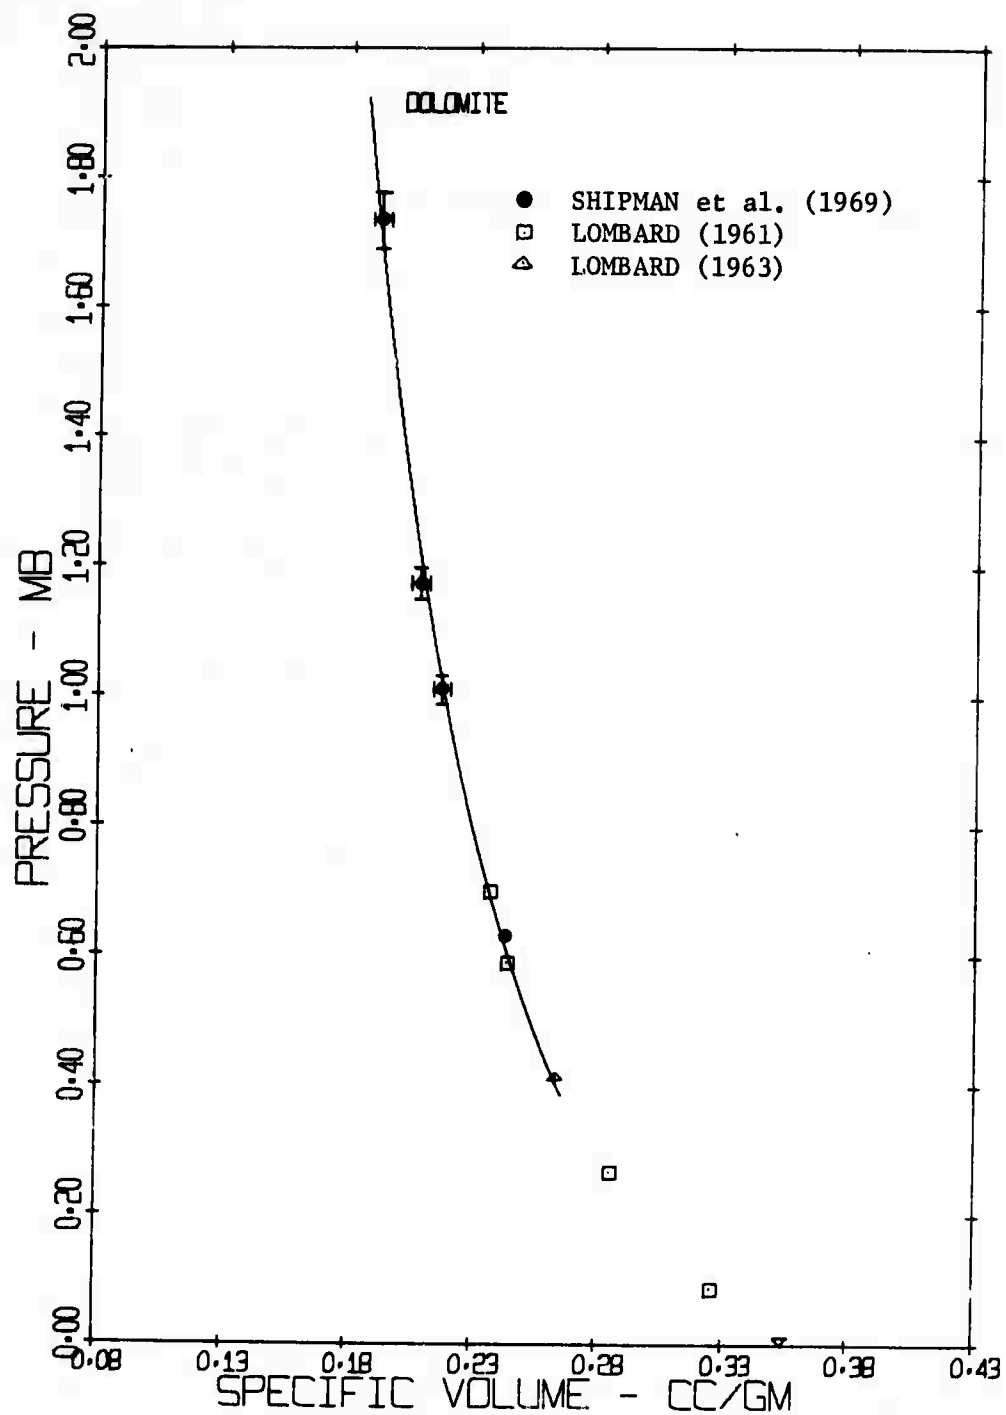
Figure 98. High-Pressure Hugoniot Data for Various Dolomites. (After Shipman et al., 1969).
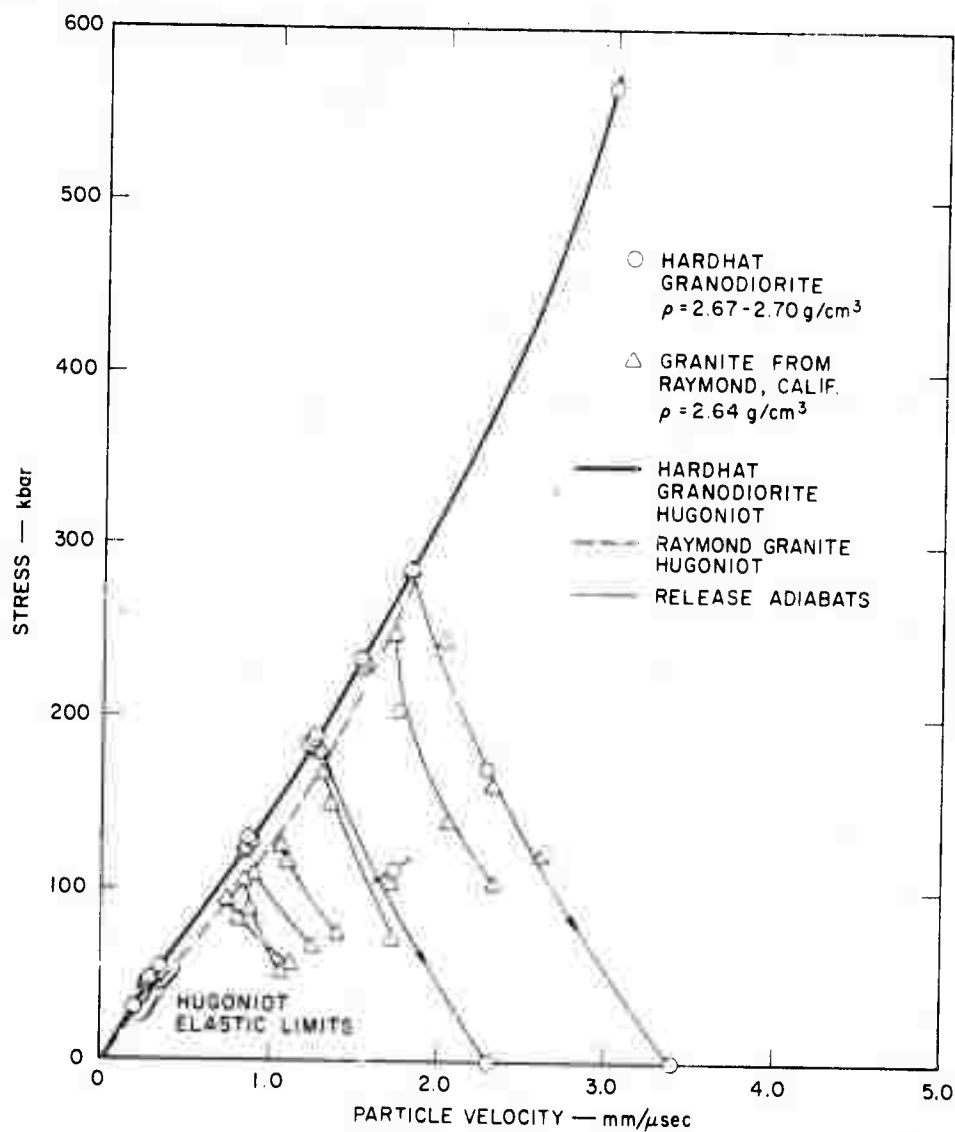
Figure 99. Hugoniot and Release Adiabat Stress-Particle Velocity
Data for Granites. Raymond granite data, Keough and
Wilkinson, 1967. Hardhat data, Petersen et al., 1968.
(After Petersen et al., 1968).

194
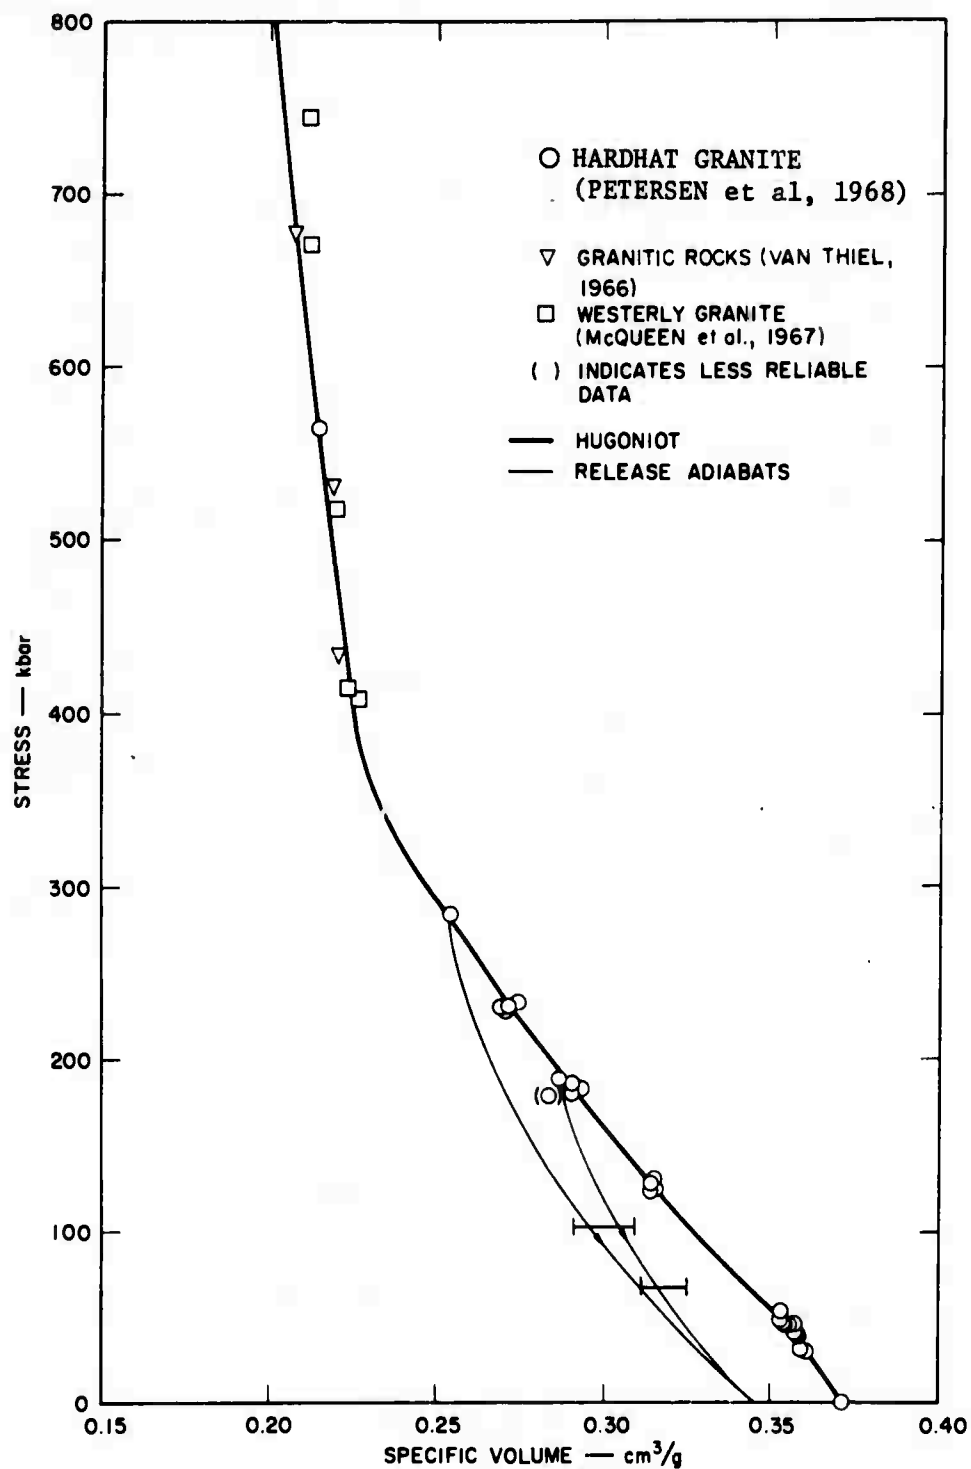
Figure 100.   Hugoniot and Release Adiabat Data for Various Granites
Calculated from Data of Figure 99.   (After Petersen et
al., 1968).
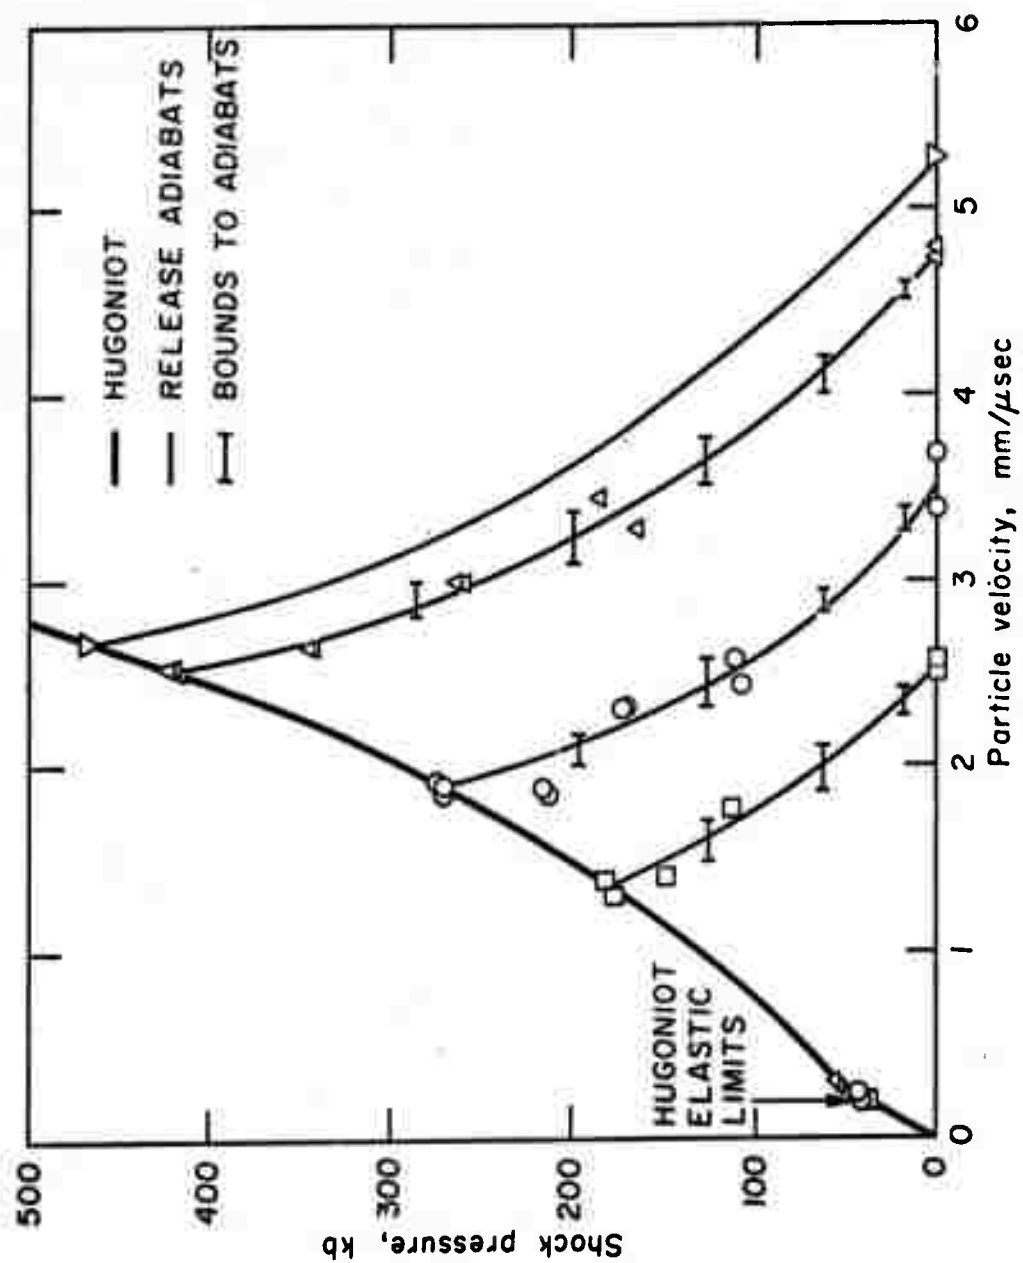
Figure 101. Hugoniot and Release Adiabat Pressure-Particle Velocity Data

196
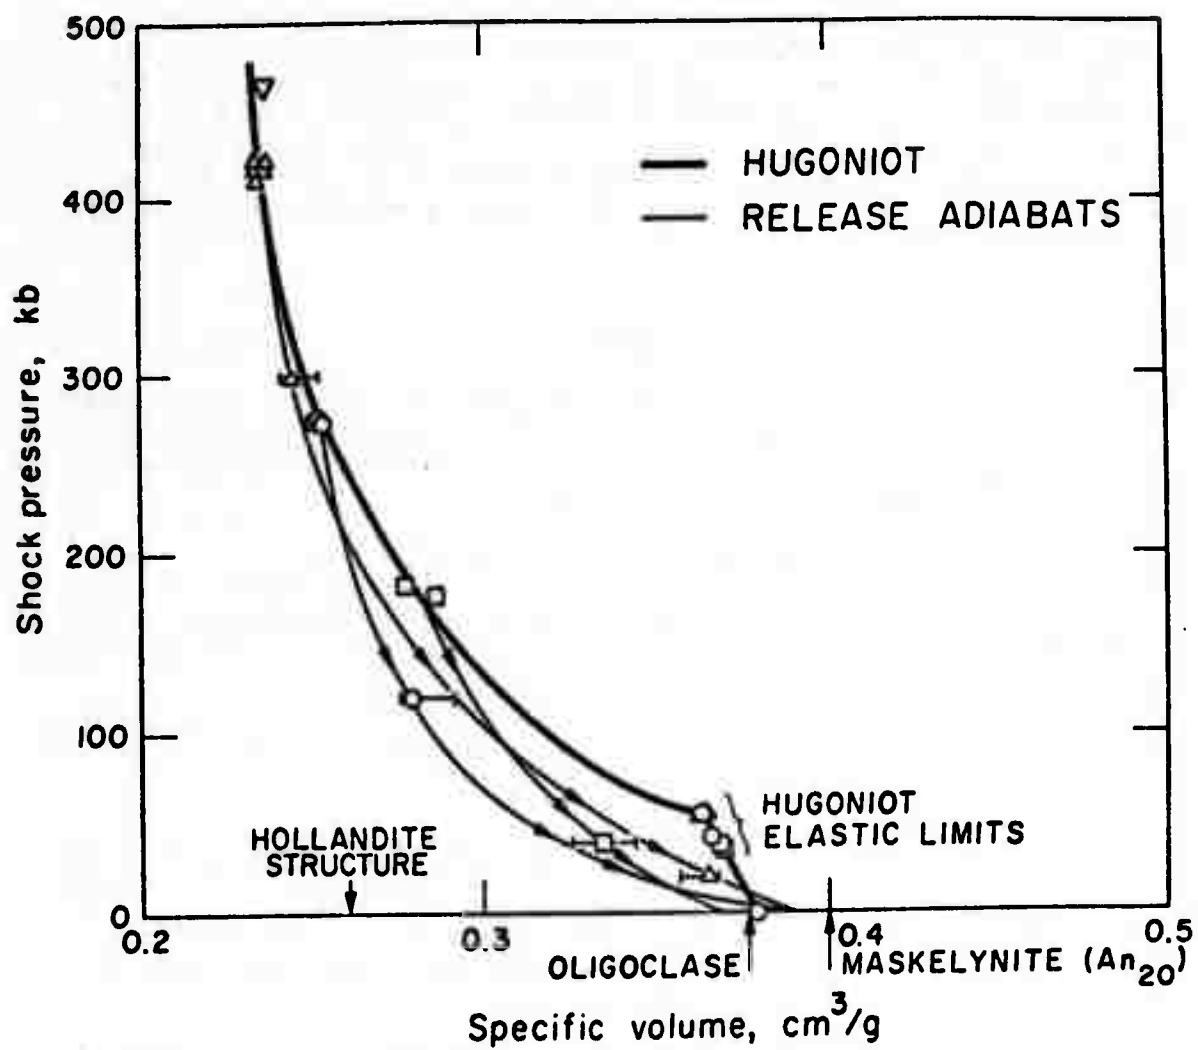
Figure 102. Hugoniot and Release Adiabat Data for Plagioclase. (After Ahrens et al., 1969b).

197

Adiabatic release from higher shock states initially occurs along steep pressure-density curves characteristic of the high-pressure phase down to pressure levels of 50 to 100 kb. Below this pressure, states are achieved in which the density is less than, or nearly equal to, the initial (unshocked) state. This surprising behavior is believed to represent the reconstructed transformation from the high-pressure phase material back to a low-pressure phase material, which in many cases has an amorphous glassy structure. Amorphous material with a zero-pressure density slightly less than single crystal, but denser than the equivalent thermal glass, is frequently obtained in samples recovered from laboratory or naturally shocked rocks. In the case of plagioclase this material is called maskelynite (Milton and DeCarli, 1963). Two general terms for this material are thetomorphic and diaplectic glass (Gibbons and Ahrens, 1971).

Summary

New data describing the rock equations of state are reviewed in the following areas: dynamic yielding of nonporous rocks, dynamic compaction of porous rocks, the effect of water on the equation of state, and very high pressure equations of state.

Dynamic yielding and hydrostatic compression data for several granites, limestone, and sandstone, indicate that immediately above the Hugoniot elastic limit these materials can retain a difference in principal stresses comparable to that existing at the dynamic yield point, the Hugoniot elastic limit. Except for Solenhofen limestone where the Hugoniot crosses the hydrostat above 20 kb, it is not yet clear at what shock stress levels strictly hydrodynamic rheological properties can be assumed.

The study of the brittle failure envelopes in principal stress coordinates of rocks under conditions of quasi-static triaxial loading has shown that failure may be predicted using explicit knowledge of all three principal stresses. It appears that a modified Drucker-Prager model can closely account for the failure surface when the yielding condition is approached by a variety of stress loading paths.

The potential application of one-dimensional stress failure tests, carried out at various strain rates, to description of material upon failure and its post-yield rheology at late times in the flow field around an underground explosion has motivated the acquisition of a body of data in this area. These data are obtained using high-speed testing machines and Hopkinson bar propagation observations. Prior to failure, marked strain-rate dependence of the dynamic (Young's) modulus is observed in volcanic tuff and granite. In contrast, no strain-rate dependent dynamic modulus is observed for Solenhofen limestone. All of the rocks tested show a moderate increase in yield strength with strain rate, as well as volume expansion as the yield point is approached.

A series of Hugoniot and release adiabat measurements for wet and dry tuff, alluvium, and slightly porous Cedar City tonalite have been carried out. Some of these data have been obtained with new techniques involving the use of a moving electrical conductor, embedded in the sample, which upon motion through a static magnetic field induces a voltage signal that is directly related to the shock particle velocity. These data, and results obtained by wave-reverberation techniques, show that the expected irreversible compaction in porous media takes place under shock compression. However, not all of the observed permanent compaction arises from a loss of porosity. Release adiabat experiments performed on fused quartz indicate a substantial portion of the density increase, which is not recovered, must take place within the silicate glass matrix of these rocks. Recovery experiments carried out on various silicate glasses substantiate this result. These data are thought to be related to the transformation of silicate glasses to stishovite-type phases at high pressure.

The effect of water content of rocks on the stress wave induced by explosions in both model and full-sclae experiments has been investigated with numerical calculations by several workers, particularly in the case of granite. At shock-stress levels greater than several kilobars, the observed stress-wave profiles can be closely matched with calculated ones by assuming that the concept of simple mixing of the adiabatic equations of state of the rock and water components is valid. In these calculations, no heat flow between components is allowed and an irreversible phase change model for the silicate component has been employed. Close description of the stress-wave profiles and peak stress attenuation with distance at lower stress levels can be obtained using either a very low (and unlikely) yield stress of ~300 bar for granite, or by incorporating into the constitutive model another mode of reducing shear stresses during later portions of the flow. A block-sliding model has recently been employed to describe this process. It should be noted that the coefficient of friction between rock surfaces, which will control this latter process, is expected to be strongly affected by water content.

Progress and knowledge of the very high pressure equations of state properties of rocks have been marked by the acquisition of a considerable body of new data, in many cases extending to two megabars, for a series of rocks and soils. The Hugoniots for silicate rocks and minerals generally show three distinct regimes: a low-pressure, a mixed-phase, and a high-pressure regime. Such behavior is observed in nearly all the materials that have been studied, with the exception of several oxides. The volume changes in the mixed-phase regime resulting from transformation to a high-pressure phase are major. They account for changes in density of from 10 to 60 percent of the zero-pressure density and produce most of the compression which silicates undergo up to 1 Mb. The new data, which includes results for granite, basalt, alluvium, tuff, limestone, and dolomite, extend the range over which the equations of state of shock-induced high pressure mineral assemblages are now known.

199

Release adiabats centered at the Hugoniot states, largely in the mixed-phase regime, have been reported for granite, plagioclase, and fused quartz. Generally these new data demonstrate that partial irreversible phase change takes place for shocks up to 200 or 300 kb, depending on the material. For adiabatic release paths centered at states at higher shock-stress levels, the release adiabats of the high-pressure phase are steep, in the pressure-density plane, down to levels of 50 to 100 kb, whereupon expansion to post-shock densities comparable to, or greater than, initial unshocked densities are observed. This striking result is believed to account for the glassy phases which are observed in quartzose and feldspar-bearing rocks in both laboratory and naturally shocked samples.

Acknowledgments

References Cited

Ahrens, T. J., and V. G. Gregson, Jr., Shock compression of crustal rocks: data for quartz, calcite and plagioclase rocks, *J. Geophys. Res.*, v. 69, p. 4833-4873, 1964.

Ahrens, T. J., D. L. Anderson, and A. E. Ringwood, Equation of state and crystal structure of high pressure phases of shocked silicates and oxides, *Rev. Geophys.*, v. 7, p. 667-707, 1969a.

Ahrens, T. J., C. F. Petersen, and J. F. Rosenberg, Shock compression of feldspars, *J. Geophys. Res.*, v. 74, p. 2727-2746, 1969b.

Ahrens, T. J., T. Takahashi, G. Davies, A proposed equation of state of stishovite, *J. Geophys. Res.*, v. 75, p. 310-316, 1970.

Anderson, Don L., and H. Kanamori, Shock-wave equations of state for rocks and minerals, *J. Geophys. Res.*, v. 20, p. 6477-6502, 1968.

Bass, R. C., H. L. Hawk, and A. J. Chabai, Hugoniot data for some geologic materials, Sandia Corporation, *SC-4903(RR)*, June 1963; see also Bass, R. C., Additional Hugoniot data for geologic materials, Sandia Corporation, *SC-RR-66-548*, October 1966.

Bjork, R. L., K. N. Kreyenhagen, and M. H. Wagner, Compressible hydrodynamic analyses of an underwater nuclear burst, *USNRDL-TRC-69-9*, CONFIDENTIAL, Shock Hydrodynamics Inc., June 1969.

Brace, W. F., Static uniaxial strain behavior of 15 rocks to 30 kb, Massachusetts Institute of Technology, Final Report, *DASA-2567*, November 1970.

Butkovich, T. R., The influence of water in rocks on underground nuclear explosion effects, Lawrence Radiation Laboratory Report, *UCRL-72558*, 1970; also *J. Geophys. Res.*, in press.

Cherry, J. T., and F. L. Petersen, Numerical simulation of stress wave propagation from underground nuclear explosions, *Proceedings of the Symposium on Engineering with Nuclear Explosions*, American Nuclear Society and U. S. Atomic Energy Commission, p. 142-220, May 1970.

Clark, S. P., Jr., unpublished data of R. G. McQueen, and H. M. Marsh, in *Handbook of Physical Constants*, edited by S. P. Clark, Jr., Memoir 97, Geological Society of America, 1966.

Davies, G., and D. L. Anderson, Revised shock wave equations of state for high pressure phases of rocks and minerals (to be published).

DiMaggio, F. L., and I. Sandler, Material models for soils, Paul Weidlinger Report to DASA, April 1970.

Dremin, A. N., S. U. Penshin, and V. F. Pogorelov, *Combust. Explos. Shock Waves*, v. 1, No. 4, p. 1, 1965.

Froula, N., The low pressure Hugoniot equation of state of anorthosite, General Motors Materials and Structures Laboratory Report to DASA, July 1968.

Gibbons, R. V., and T. J. Ahrens, Shock metamorphism of silicate glasses, *J. Geophys. Res.*, in press, 1971.

Green, S. J., and R. D. Perkins, Uniaxial compression tests at strain rates from 10^{-4}/sec to 10^4/sec on three geological materials, General Motors Materials and Structures Laboratory Report *MSL68-6*, April 1968.

Jones, A. A., W. M. Isbell, F. H. Shipman, R. D. Perkins, S. J. Green, and C. J. Maiden, Material property measurements for selected materials, in General Motors Materials and Structures Laboratory Report, *NAS2-3427*, 1968.

Jones, A. H., and N. A. Froula, Uniaxial strain behavior of four geological materials to 50 kilobars, General Motors Materials and Structures Laboratory Report, *DASA 2209, MSL-68-19*, March 1969.

La Mori, P. N., Compressibility of three rocks, (A) Westerly granite and Solenhofen limestone to 40 kbar and 300°C, (B) Cedar City tonalite to 40 kbar at room temperature, Battelle Memorial Institute Report, *DASA 2151*, August 1968.

Lombard, D. B., Rock Equations of State, University of California, Lawrence Radiation Laboratory Memorandum *UOPK 63-2*, January 1963.

Lombard, D. B., The Hugoniot equation of state of rocks, *UCRL-6311*, February 1961.

Lysne, P. C., A comparison of calculated and measured low-stress Hugoniots and release adiabats of dry and water-saturated tuff, *J. Geophys. Res.*, v. 75, p. 4375-4386, 1970.

Milton, D. J., and P. S. DeCarli, Maskelynite: formation by explosive shock, *Science*, v. 140, p. 670-671, 1963.

Mogi, K., Effect of the intermediate principal stress on rock failure, *J. Geophys. Res.*, v. 72, p. 5117-5132, 1967.

Murri, W. J., and G. D. Anderson, Hugoniot elastic limit of single-crystal sodium chloride, *J. Appl. Phys.*, v. 41, p. 3521-3525, 1970.

McKay, M. W., and C. S. Godfrey, Study of spherically diverging shock waves in earth media, Physics International Report *PIFR-064*, *DASA-2223*, March 1969.

McKay, M. W., R. R. Shadel, R. P. Swift, and C. Young, Experimental studies of stress-wave propagation in earth media, Physics International Report, *DASA 2525*, Sept. 1970.

McQueen, R. G., S. P. Marsh, and J. N. Fritz, Hugoniot equation of state of twelve rocks, *J. Geophys. Res.*, v. 72, p. 4999, 1967.

McQueen, R. G., Shock-wave data and equation of state, in *Seismic Coupling*, edited by G. Simmons, p. 53-106, VESIAC Report, Geophysics Laboratory, University of Michigan, 1968.

Perkins, R. D., and S. J. Green, Uniaxial stress behavior of poryphyritic tonalite at strain rates to 10^3/second, General Motors Materials and Structures Laboratory Report *MSL-68-42* for DASA, October 1968.

Peselnick, L., Elastic constants of Solenhofen limestone and their dependence on density and saturation, *J. Geophys. Res.*, v. 67, p. 4441-4448, 1962.

Petersen, C. F., W. J. Murri, G. D. Anderson, and C. F. Allen, Equation of state of rocks, Stanford Research Institute, Interim Technical Report to Atomic Energy Commission, Lawrence Radiation Laboratory, June 1968.

Petersen, C. F., W. J. Murri, and R. W. Gates, Dynamic Properties of Rocks, Stanford Research Institute Report, *DASA 2298*, June 1969; see also Petersen, C. F., W. J. Murri, and M. Cowperthwaite, Hugoniot and release-adiabat measurements for selected geologic materials, *J. Geophys. Res.*, v. 15, p. 2063-2072, 1970.

Riney, T. D., S. K. Garg, J. W. Kirsch, L. W. Morland, and C. R. Hastings, Stress wave effects in inhomogeneous and porous earth materials, Systems, Science and Software Report, *DASA 2498*, March 1970.

Rosenberg, J. T., T. J. Ahrens, and C. F. Petersen, Dynamic properties of rocks, Stanford Research Institute Report, *DASA 2112*, July 1968.

Shipman, F. H., W. M. Isbell, and A. H. Jones, High pressure Hugoniot measurements for several Nevada Test Site rocks, General Motors Materials and Structures Laboratory, *DASA-2214, MSL-68-15*, March 1969.

Simmons, G., editor, Seismic Coupling, VESIAC Report, University of Michigan, 1968.

Stephens, D. R., The hydrostatic compression of eight rocks, *J. Geophys. Res.*, v. 69, p. 2967, 1964.

203

Stephens, D. R., E. M. Lilley, and H. Louis, Pressure-volume equation of state of consolidated and fractured rocks to 40 kb, *Int. J. Rock Mech. Min. Sci.*, v. 7, p. 257, 1970.

Swanson, S. R., Development of constitutive equations for rocks, University of Utah, Ph.D. Thesis, 1969.

van Thiel, M., Compendium of shock wave data, University of California, Radiation Laboratory Report, *UCRL50108*, 1966.

Wackerle, J., Shock-wave compression of quartz, *J. Appl. Phys.*, v. 33, p. 922, 1962.

Wagner, M. H., and N. A. Louie, Hardhat/Piledriver ground motion calculations, Shock Hydrodynamics Report, *SAMSO-TR-69-47*, March 1967.

Wang, C., Phase transitions in rocks under shock compression, *Earth Planet. Sci. Lett.*, v. 3, p. 107, 1967.

Wiedermann, A. H., and O. E. Curth, Shock unloading characteristics of crushable rocks, *WL-TDR-64-52*, Contract AF29(601)-6007, June 1964.

204

DISCUSSION OF EQUATIONS OF STATE

MR. TRULIO: Are release adiabat data available for rocks with different water contents? For example, simple equation of state calculations indicate that a shock of 50 to 70 kb will deliver enough energy to water to initiate boiling upon release.

MR. AHRENS: There are some data for water-containing materials, such as alluvium and tuff. Several years ago Anderson measured release adiabats for frozen materials which I believe are pertinent.

I am not aware of any data that pinpoints the region along the Hugoniot where initial vaporization takes place. There have been experiments conducted in which effects, such as very high free surface velocities, were noted. These were inferred to be the result of boiling. However, a careful study of these phenomena has not been reported.

MR. TRULIO: The stress at which boiling will initiate is important, because stresses of 50 to 70 kb are never reached by much material, for example from a surface burst, and it is a question of whether or not it is important to know details of volatile component behavior.

MR. STEPHENS: The maximum pressure to which a water-containing rock must be shocked to begin to enter the two-phase region upon loading of course varies with the rock, the water content, and especially the amount of gas-filled porosity--porosity which is not filled with water. For some tuffs, water in the rock will begin to boil upon unloading from peak pressures of 100 kb or less.

On the other hand, at lower pressures, such as 20 kb, although the water in the rock will not boil upon unloading, the volume expansion due to water at low pressure is appreciable, due to residual heat. This may not be a small effect when using a release adiabat in a calculation.

MR. CHERRY: I believe what is needed now is the release states from a saturated material as, for example, it is shocked to 100 kb and released. After those data are available, I believe the models can be much better defined. Right now your answers depend on what model you use, whether you assume that the pressure of the water and the pressure of the rock equilibrate at the shock front, or if they follow their individual isentropes and the release state then is determined simply by adding the isentropes of the rock and the isentrope of the water. I think you can certainly construct a model that will show appreciable effects at peak shock pressures corresponding to 70 kb and above. Whether it is really true or not will not be known until release data are available.

MR. GODFREY: I am concerned that the release data that will become available take place in about a microsecond while the data we really need are data on release times of the order of 100 ms. That is five

orders of magnitude slower in time. I think it is true that rocks which have gone through phase transformations at high pressure have in general transformed back to their original state by the time they are recovered. So, we do not know the time scale in which the transition takes place. We await the release data with great expectation, but it may not be relevant.

MR. CHERRY: It is possible that the material strength of a rock relaxes behind the Hugoniot elastic limit so that the hydrostat and Hugoniot data should actually correspond and that the steep release path you measure represents the reoccurrence of strength during the release phase of the experiment?

MR. AHRENS: It is not clear what the stress differences are in the high-pressure regime well above where one can compare the hydrostat and Hugoniot. I think this is a region that is important to the calculations and where equation of state data are needed. The release, for example, from fused quartz, which is thought to be a relatively weak material even on compression, indicates irreversible compaction of the crystalline material, and suggests that the irreversible compaction is not due to an elastoplastic effect, but is a thermodynamic effect due to a partial phase change. For granite, clearly both effects could occur. I don't believe the state of the art in equation of state measurements is at a point now where we really can tell the difference. We do not know the stress deviators in granite at 200 kb. For some materials we know that hysteretic behavior is due in part to phase changes, but it is not clear that it may not have an elastoplastic component.

MR. GODFREY: Just intuitively you would expect that during a phase change the newly formed material would not remember what stress state it was in before it transformed. It is hard to see how these stresses can be maintained through a phase change.

MR. MADDEN: When you consider the time scale involved here, I wonder if some other physical processes are not more important for seismic effects. For instance, in a time scale of a fraction of a second, how about the outflow of gas due to the permeability of the rock and the effect this has on the pressure? The question is what are the physical phenomena that take place long after the phase changes we have discussed, say in a tenth of a second. These may be important for the generation of the low frequency components of the seismic signal.

MR. STEPHENS: In the scale of a tenth of a second, permeability would not be particularly important.

MR. MADDEN: You mean to say that at a tenth of a second there is no chance for the movement outwards of the confined gases?

MR. STEPHENS: Or water. No, I don't think so.

MR. RINEY: It might be of interest to point out that calculations obtained with the normal computer codes do not account for any possibility of diffusion. It could be that gross diffusion plays no role, but relative motion of the fluid in the pores could have a significant effect on the signal, I believe. We are involved in developing a theory which does account for this relative motion of the fluid in the matrix in wave propagation studies.

MR. STEPHENS: Do you calculate water or gas motion in the pores in this time frame?

MR. RINEY: You could have relative motion in the pores in that time. I am not saying fluids move from one pore to another, but there is relative motion. You can consider it, if you want, just a very idealized situation where you have a wave propagating in a direction parallel, say, of water in a tuff, and you would very definitely get a tremendous flow differential velocity in that direction.

MR. ROTENBERG: I would like to question whether any of this is very important. I think it is unfortunate that we had to miss the seismology session yesterday. If we talk in terms of teleseismic distances at frequences of the order of a cycle, then we seem to have some evidence that all of these details are unimportant. That is, the signals we observe at teleseismic distances in the frequency range of one to three cycles are insensitive to what is going on at the site of the explosion. Once you are given the motion at the wall surface of a cavity, then all of the environment, all of the equations of state do not seem to make very much difference at large distances.

MR. CHERRY: I don't know about large distances, but at LRL we continuously monitor the activities at the test site at a distance of 250 km. We observe enormous variations in signal content of the record as we go from alluvium to granite to tuff. In fact, we can even see differences as we approach the water table in alluvium as the shots are deeper. I think it might be interesting to try to clarify those seeming discrepancies between the 250 km recordings we have with the teleseismic magnitudes that were quoted yesterday.

MR. EVERNDEN: Before this goes too far, we had better clarify what I did intend to say. I did not say that signals in alluvium looked like those in granite. I did not say that you cannot observe an effect as shot depths approach the water table in alluvium. The signal amplitude of a shot below the water table is two orders of magnitude higher than for a shot above the water table. These are for low yields.

What I said was that we could not distinguish Mesa tuff from granite and salt. For 100 tons yield you cannot differentiate alluvium signals from granite signals in amplitude. But at above a kiloton there is a real difference between alluvium and what you observe in other materials.

Valley tuff will have a different seismic magnitude for a given yield than will granite or will the Mesa tuff for yields above 4 kt. Below that, its curve appears on the margin of the Mesa tuff and the granite curve. Of course, water-saturated materials have a response close to pure water. That is, for saturated rocks such as alluvium, saturated Mesa tuff. and saturated valley tuff, the seismic magnitudes are very near to water.

MR. SILLS: . I think the value of the codes is inside of the 50-kb region, and I think there is a real value to codes in this area showing coupling and decoupling as part of the motion of the wall. This is the type of thing you are concerned with.

MR. ROTENBERG: I specifically said once you get to the wall. I am not discounting the value of codes up to the wall.

MR. GODFREY: There is a paper by Higgins and Buktovich of LRL that I think is relevant to this discussion. I will just mention what the paper was about, and what the conclusions were. They attempted to correlate the radii of all of the underground nuclear shots with over-burden pressure, the strength of the material, and water content. They found the correlation without factoring in the strength of the material was just as good as trying to include the strength. In other words, they found a good correlation by assuming that granite was no stronger than alluvium in so far as its resistance to expansion and the forma-tion of a cavity, which was rather startling. On the other hand, we might not be startled considering that these are all jointed media; perhaps they are only as strong as the joints.

I am bringing this up because it may be that considerations of strength as far as this problem, ARPA's problem in this area, is not terribly relevant.

As far as Hugoniots go, we did some studies where we compared one megaton in nonporous limestone with a megaton in nonporous granite, for example, assuming they were both weak and had the same effective strength. You could hardly tell the difference in the profiles. Most of the silicate rock Hugoniots looked very much alike, so I question really whether in nonporous media, whether the detailed differences in the Hugoniots make much difference. Porosity is a different thing.

MR. STEPHENS: In regard to the Higgins and Buktovich correlation of cavity radius, what you say is certainly true; on the other hand the standard deviation in the radii which they calculated with their corre-lation was 15 percent, and the question then arises whether this is a satisfactory enough determination of the cavity radius, which as I understand is one of the things that taking the strength and equation of state into account, one can compute quite precisely.

MR. FRASIER: I wish to ask a question about dispersing mechanisms that were suggested yesterday and mentioned today. The reason I ask this is

208

that at what we call high frequencies (from 1 to 5 cps) the body waves from explosions and earthquakes do not really show measurable dispersion. Part of the reason for this is that the unknown geologic factor is a source of attenuation of the earth at teleseismic distances. I would like to know what frequencies you are really talking about in terms of measuring effects of a close-in source, and whether these would be measurable in the one cycle range at all?

MR. RINEY: I really at this time can't say what kind of dispersion you will see, but we will try to calculate it.

MR. FRASIER: Could that be seen at close-in stations?

MR. RINEY: Well, it depends. Of course, your media are changing, so most of these calculations are with one type, making it quite difficult to go on from that step and apply it to a real-life situation where your geological structure may be changing. I really can't say.

MR. TRULIO: But even for small-yield shots, the gages that are placed in the field have maximum response of about 10 kc. So you are talking about very high frequencies, when you talk about megacycles. But if you are thinking of scales of distances like the pore size, you have to talk about very high frequencies.

MR. AHRENS: Commenting on Dr. Frasier's question, it is my understanding that attenuation of the effect he is referring to is for teleseismic body waves which have gone deep into the mantle and returned to the surface again. I would think that except for possibly the small propagation path length through the crust, and possibly through the low velocity zone in the upper mantle, most of the path would be fluid free; so you would not expect a viscosity dependent dispersion to play an important role for these wave paths.

MR. GRINE: I would like to make a comment on measuring release states. The release states are inherently more difficult to measure than the Hugoniot, particularly at high stresses and particularly in porous medium, because the gages must survive through the whole crushing phase and keep on recording during release. In a porous medium with water, you have different pieces moving at different velocities, and different pieces of your gage also move at different velocities, and the gage really does not last very long. The higher the stresses, of course, the bigger these differential velocities are and the shorter the recording time.

We are trying a variety of techniques with gages, making gages thicker and bigger and so on, and we do measure release states. Although we can't get to the millisecond range, which Chuck Godfrey says they would really like to see, we can at least measure in times from a tenth of a microsecond to a few microseconds, and see if we can see rate-dependent effects in inverse phase changes, vaporization, and so on over that time scale. That is the best we can do right now.

COL. RUSSELL: We put together this seminar to give those involved in code calculations an opportunity to informally exchange information on the ILLIAC. Hopefully we are going to conduct this thing at three different levels. First I have asked Dave McIntyre from the University of Illinois ILLIAC group to talk in some more sophisticated detail about the hardware and the software and the programming work that they have done at Illinois on the ILLIAC. This will give you a feel for where they stand now, what the problem areas are, and where you can go to get more information on the machine.

Then I want to talk very briefly and in general terms about ARPA's plans and other plans to reconfigure some codes into the ILLIAC language, and then finally I would solicit your comments on possible areas, problems, or codes that you might think are appropriate to be attacked on the ILLIAC.

We will start off with Dave.

MR. MC INTYRE: I would like to start by going into more detail on the processing element, which is the basic building block of the system. If you remember from yesterday, there are 64 of those, and they are all driven by a single instruction stream. The processing element is basically a four-register computer, similar to the old 7094 in that it has something like an accumulator and M/Q register; an S register, which is used to store intermediate results; and an R register, which participates in the routing operation. All of these registers are 64 bits wide.

There is an X register, which is 18 bits wide and is used to modify base addresses. This is what allows you to reference different memory locations in the memory of different processing elements. The memory is 2,048 words. It is a semiconductor memory with an access time of about 200 nsec.

In order to access different words in different processing elements, you load the X register with different numbers. You might load it with five in processing element No. 1 and ten in processing element No. 2. When the control unit sends down a command "fetch from location zero indexed by the X register," you would fetch from location No. 5 in processing element No. 1 and from location No. 10 in processing Element No. 2.

The R register in processing element i is wired to the R register in processing element i plus 1, and in processing element i minus 1, and in the 64 processing element, the R register is hard wired to the R register in processing element 1. You can think of the routing operation which distributes operands among the processing elements as essentially a shift on a very long register, and the shift is in the route.

Preceding page blank

In addition to being wired to the neighboring processing element, the R register in the i^{th}PE is wired to the R register in a PE that is located eight away from where it is. So if you have to do very long distance routing you do it in jumps of eight rather than in jumps of one. The routing operation is very fast. It requires two clocks. A clock is 60 nsec.

The hardware automatically decodes an arbitrary distance route into multiples of eight and one. A route of distance 20 would be two routes of distance eight and four routes of distance one.

If you want to multiply, you load the two operands in register A and the M/Q register, which we call register B. You say multiply, and the result comes back in register A.

In addition to these registers, there is a series of one-bit registers in which you can store logical results, and two of the one-bit registers tell the processing element if it is on or off. You can do such things as transfer a bit from the A register into one of these one-bit registers and turn the processing element off. That register is called the mode register.

Are there questions on the processing elements?

QUESTION: What was the low fetch time from the memory to R?

MR. MC INTYRE: The fetch time is seven clocks, or 420 nsec. The access time at the memory is only 200 nsec, but unfortunately this is a semiconductor memory. The PE is also built out of semiconductor components, but a different family, so there are voltage differences. We have to go through an interface, which slows us down. The fetch time is seven clocks, and store time is six clocks.

QUESTION: You can fetch to R as easily as to A and B?

MR. MC INTYRE: Yes. Let me make a few remarks on the control unit. In it you have a fairly large instruction stack of 64 64-bit words. Each 64-bit word can store two instructions, so you can get up to 128 instructions in a control unit. These are divided into eight word blocks. As you are executing down this instruction stream, which originates in the processing element memories and is stored across the processing element memories, and when you have executed the fourth instruction word in an eight word block, the hardware looks to see if the next eight words are in this register file, this program stack. If they are, nothing is done. If they are not, the hardware initiates a fetch to bring the next eight words in. By using that kind of simple strategy, it turns out you are very seldom held up waiting for the control unit to fetch instructions.

Also in the control unit there is a local data buffer, which is 64 64-bit words. There is a fixed-point arithmetic unit to do

212

logical operations and simple fixed-point arithmetic, and there is also a queue consisting of eight instructions which feeds the microsequence generator, which in turn drives the array of processing elements.

The object of the game is to allow the control unit to process instructions faster than the PE's can execute them, and to fill this queue up, so that the processing elements are always kept busy.

The fact that there are 64 processing elements and the word size is 64 bits is not just because both are powers of two. The control unit has to make a decision. All branches in the instruction stream are performed by the control unit, but occasionally it has to check on the status of what is going on in the array. For instance, it might be nice if, while you were doing a hydro calculation, you knew when one of the processing units computed a mass which was perhaps negative. To get that kind of information the control unit has to copy one bit from each of the processing elements. It copies them off into a 64-bit word, and then, scans the 64-bit word to see if it has all zeros, which says all masses are positive, or if it happened to run into a one, which says somebody computed a negative mass. Then it can branch based on that information.

There is another section of the control unit, called the memory service unit, which coordinates the requests on the memory and resolves conflicts. There are several units making demands on the memory. The control unit makes demands to fetch the instruction stream and load the local data buffer. The processing elements copy operands into their operating registers. And the I/O system makes demands. The memory service unit resolves those conflicts and assigns the I/O system the lowest priority in order to get the memory.

The backup to the 131-K memory is a 10^9-bit disc, which is a rotating device. If you are performing I/O from the disc, and the memory is being used heavily by the processing elements, you may run into the situation whereby the I/O cannot get in to use the memories. The disc is turning, so you may lose the address on the disc. If this occurs, the memory service unit allows the I/O to have top priority, and I/O steals the memory cycle and gets in.

QUESTION: What is the physical size of that disc?

MR. MC INTYRE: The disc is 36 in. in diameter. It is a head per track device. There are actually 13 storage units or 13 discs. We read off both sides of the disc, 128 tracks at a time, and that allows us to get the big transfer rate. You read maybe 10 or 11 tracks at once on a conventional disc.

Are there any other questions on the hardware?

QUESTION: Is the compiler going to treat this disc as a separate device, or will this automatically be blended into the operational procedure?

MR. MC INTYRE: No, it treats it as a separate device and all I/O has to be explicitly stated, at least in the compilers we are working on right now. It would be possible to develop a compiler that did implicit I/O, but I am afraid you would pay some overhead for that.

QUESTION: Will you have some analogous statements like the 6600 has, called buffer in and buffer out, that will treat that disc to the processing element's memory, or at least I/O macros?

MR. MC INTYRE: Virtually all input and output is buffered input and output, in that sense. The I/O subsystem returns the status work to the control unit, saying either I have completed the I/O transaction, or it is still in progress, or I have an I/O fault.

QUESTION: At present, this will be handled by macros in the assembly language? Today?

MR. MC INTYRE: Today, yes.

MR. RANDALL: Some of it.

MR. MC INTYRE: I think that Mike will go over an I/O transaction when he talks about software.

QUESTION: Do you envision that, in a difficult hydro calculation, it would be possible to put part of the mesh or the bulk of the mesh on the disc?

MR. MC INTYRE: I think it is important that you be able to, because the memory is actually of a fairly modest size, and with that large a commutation power, you exhaust the operands which you can hold in that memory in a very short time.

QUESTION: The total access time is what?

MR. MC INTYRE: The disc rotates once every 40 msec, so the average length would be 20 msec. But there is an interesting piece of hardware that is associated with it which is an I/O request querer. If you can stack several I/O descriptors into that querer, the hardware is automatically reading the address on the disc that is passing under the head, and it will initiate the one which minimizes latency. So if you can gang up several I/O requests instead of seeing an average of perhaps 20 msec latency, you may see 10 or even less.

We did some calculations based on the SHELL code out of the weapons lab, and it turned out, that for the 2-D problem, we were very close to being I/O bound, within the noise of the calculation. It did not really matter. For the 3-D problem, we were a little bit I/O bound. If the I/O system had been twice as fast, we would not have been I/O bound. What I am saying is that we could compute faster than we could bring operands in and put them out.

214

QUESTION: Did you get the 3-D calculation on that disc?

MR. MC INTYRE: Yes.

QUESTION: What was the grid size?

MR. MC INTYRE: I think something like 100 by 100 by 100.

QUESTION: So that is 10^6 times 50 variables?

MR. MC INTYRE: Fifty variables? What kind of hydro is that?

QUESTION: I am sorry. It was plastoelastic.

MR. MC INTYRE: In the next 10 or 15 min, I would like to make a couple of comments about how you use the architecture to do two-dimensional problems. I talked about one-dimensional problems yesterday. Then we will have a talk on the details of the software.

There are two difficulties in using ILLIAC IV. One is to find some way to use the simultaneous computation capability. The other is distributing the operands among those discrete memories in such a way that every processing element can latch onto any piece of information that it needs for the calculation.

If you are working with vectors, and the vectors are of arbitrary length, generally the way you store them in the processing element memories--let this be memory 64, 63, 1 and 2--is to just start out at a location--call it ψ--and store U-1, there, U-2 at the same location in processing element 2's memory, U-63, U-64, and then just wrap around and put U-65, U-65, U-66, and so forth. This preserves essentially the connectedness of the vector, in the sense that the left and the right neighbors are processing element memories which are close to PE-1. By close, I mean that they can be gotten in very quick routes of distance 1 or 2 or 3.

You occasionally waste some memory because the vector is not of a length that is a multiple of 64. You are going to have to store it at a level, say, $\psi + 3$, where there are no components to fit. You can either pad it out with dummy components, or just throw away that memory space.

When you are working in two dimensions, you need some way to store a matrix, and there are a couple of schemes for doing that. Let me consider ILLIAC IV as composed of four processing elements. This is their memory proceeding down this way. If you have a four by four matrix, you store U-11 in processing element 1's memory at a given location, say ψ, and U-12, U-14. You store U-21 in the second row at $\psi + 1$, and so forth.

215

If you observe, this is PE-3. If you are doing finite-difference calculations, you are generally working on something like a five-point star where values at this mesh point are dependent upon values at this mesh point, and four neighboring values. Occasionally you extend this to these. You will see with this kind of storage, when you want to operate on U-33, you have readily accessible in processing element three the north neighbor and the southern neighbor. Then in the neighboring PE, PE-2, you have available the east neighbor and the west neighbor. So that those can be obtained using a route of just distance 1. Is that clear?

Now, if you are doing something like the cycle on a hydro integration, you may want to do certain things on the interior of the mesh but something very different on the boundaries. On this kind of storage, you can access in parallel all of the values on the top boundary and the bottom boundary because they are each stored in a different processing element's memory. You can copy them to the operating registers in parallel, and then you can adjust your boundary values.

But if you look down this boundary, or down this one, you see that they are all stored in a single processing element. If you want to adjust values at these mesh points, you have to do it sequentially with this kind of storage. Sometimes it is acceptable, depending upon how quick the operation is, to just do it sequentially. Occasionally, and more particularly in matrix computations, you would like to be able to access both rows and columns at the same time, or with equal ease, that is, you would like to be able to access in parallel rows and columns.

There is a method of storage called skewed storage which allows you to do that. In skewed storage you start out storing the matrix much in the same fashion as in straight storage, but then you rotate it the distance one PE to the right and store U-21 in PE No. 2's memory, and you wrap around. Then you repeat the process. This is processing element 1, 2, 3, 4. You can see once again the first row is stored in separate processing element memories, so you can access them in parallel. The first row, for example is stored in separate processing element memories. But so is the first column with this kind of storage, because the first column lies here, and to access the first column simultaneously all you do is load the X register and PE-1 with zero, PE-2 with 1, PE-3 with 2, and PE-4 with 3. And you say fetch, with the location ψ indexed by the contents of the X register, and that does the memory operation in PE-1 at this location, PE-2, and PE-4 here. You have the first column, so you can adjust those values in parallel.

If you want to get the second column, you just rotate this index pattern around one, and do the same thing.

QUESTION: Are there some instructions to do that index arithmetic like 64 wrap around to set these index registers up to get the J columns?

216

MR. MC INTYRE: Originally they are loaded by the compiler just at object time, and you just have to manipulate them, either arithmetically or by logical operations.

QUESTION: Would it be a table of 64 different index register values?

MR. MC INTYRE: No, you just distribute them.

QUESTION: But you have to wrap around.

MR. MC INTYRE: Yes.

QUESTION: So it is about three or four instructions to set up your index register.

MR. MC INTYRE: Something like that--fetch to R, route, a distance, load into X--three or four.

Now, if we are doing three-dimensional calculations, the easiest way to think about it is to put three two-dimensional planes in the memory. You can't core contain any meaningful three-dimensional calculation anyway. So you might as well just have three two-dimensional planes, and bring one in while you are processing one, while you are writing one out.

I guess I should summarize what the machine does well and not so well. The machine does finite-difference calculations or mesh calculations very well, if you will accept meshes that are multiples of 64. You often can get efficiencies in excess of 80 or 85 percent. By efficiencies, I mean the average number of processing elements turned on during the calculation is approximately 80 percent of 64. If you want to have arbitrary size meshes, you sometimes suffer a little in your efficiencies, but you seldom degrade below 60 percent.

Explicit finite-difference calculations are probably easiest, but implicit ones work, too. You have only to solve nonlinear equations in the implicit calculations, which is done either by linearizing or by using successive substitutions.

Matrix calculations go very well on the machine, and efficiencies there are generally in excess of 50 percent.

Table look-up problems, if the table is relatively small, go fairly efficiently. But when the table is very large and can't be contained in a single processing element's memory, they go very poorly.

Particle-moving problems go from modest to very poor, depending upon the type of problem. Particles in cell hydrodynamics generally shows modest performance of perhaps over 50 percent.

217

Nonlinear radiation transport, where the particles affect the absorption properties of the medium through which they are being transferred, goes very poorly. One might expect from 40 percent to 25 percent on those kinds of calculations.

Are there questions?

QUESTION: You gave a number on PIC.

MR. MC INTYRE: Yes, we did a study with Los Alamos on a PIC plasma code, and it turned out there that the efficiency during the particle-moving phase was about 80 percent. In some cycles, it dropped as low as 60 percent. The distribution in the efficiencies looked something like this, where the cutoff here was about 60 percent and here about 80 percent.

In all fairness, the tail did not go out that far. Different time cycles resulted in different efficiencies based on the distribution of the particles, because the particles migrated with the calculation, crossing processing element boundaries.

QUESTION: We have been imagining that that would be a substantial disadvantage.

MR. MC INTYRE: The stability criteria under which they were operating the code restricted it so that particles could only cross from one cell to a neighboring cell. I am not so sure that is really a sound stability criterion, and if you wanted to allow particles to cross many cells, you have a difficult problem in programming ILLIAC IV.

QUESTION: How much are the efficiencies related to the skill with which the programming is made?

MR. MC INTYRE: It very definitely depends upon the skill. As some people in the audience can tell you, you can spend considerable time formulating these problems in an optimum way. It may turn out that you probably can come up with a method of adapting your problem to this architecture fairly quickly, but then you start to ask yourself is there a better way? You can continue refining like that for a considerable time.

QUESTION: The second part of this question is this: If you continue refining for a considerable time, does this cost you $1500 an hour, or is there some sort of a simulation program that operates on another machine?

MR. MC INTYRE: Right now there is a bit-by-bit simulator that runs on the Burroughs 5500. You can check out your codes on it, but it is very slow. It is about a million times slower than ILLIAC IV, so you can't run many cycles with your calculation. If you are playing around or have written a code for the machine using a couple of different

approaches, it will cost you $1500 an hour before you even settle on a production code.

QUESTION: When you say you have done these, have you actually done them on the ILLIAC, or just simulated them?

MR. MC INTYRE: Just simulated them. Actually, what we did on the Los Alamos code was come up with a memory allocation scheme and take snapshots of the distribution of particles in a running code, or from a running code. Then, based on those snapshots and the distribution of particles, we calculated the efficiencies that the code would achieve. We did not actually move the particles in the code, because it would have taken too long.

QUESTION: Don't you have a timing simulator, which does not do the calculation, but tells you the efficiency of the number of PE's that will be used?

MR. MC INTYRE: We have a pseudo-timing simulator, but it is rather difficult to estimate the time correctly because of the overlap and concurrency between the control unit and the processing elements, and because certain operations overlap in the processing elements. It is very difficult to time the machine in any way other than just to execute the code.

QUESTION: Especially with those conditionals coming in.

MR. MC INTYRE: Right. You don't know how many you will have.

QUESTION: But pending conditionals, does that timing simulator give you ...?

MR. MC INTYRE: It counts the clocks on the instruction stream.

QUESTION: So this tuning could be done on that timing simulator.

MR. MC INTYRE: Right.

QUESTION: Trying to get the efficiency, which would be the number of PE's that would be concurrently operating?

MR. MC INTYRE: Yes.

QUESTION: Did you say yesterday that this machine was going to be down in its operation or from its operational mode every 5 hr?

MR. MC INTYRE: Oh, no, I didn't say that at all. I said the average time to fail was 5 hr. It turns out the average time to repair and verify is 30 min. Availability is something like 90 percent, which is not bad for a big machine. You see, we have a series of programs which detect hardware failures, and these are run intermittently. Then, when

219

an error is detected, we have a series of programs which isolate the processing element where it occurred.

We are building 70 processing elements so we have six spares. What you do is to unplug one of those processing units, swing it out, and plug a new one in. Then you run a verification, which requires about 5 min. Then you take the bad one back, locate the failed part, and repair it.

QUESTION: I am not too used to that type of reliability on computers. Suppose you are running a hydro code that takes an hour, or maybe even 20 min. Does the machine have enough parity checks to guarantee that when machine failure occurs we won't get catastrophic answers?

MR. MC INTYRE: No, the machine has no parity checks at any point. As a matter of fact, the only way you can determine errors is to run a confidence diagnostics program which exercises all of the branches in the logic in the PE. You can compute 64 answers at the same time, so if one fellow gets a different answer, you would probably suspect that there is a logic error there.

QUESTION: Yes, but if I am running a hydro code for 20 min, and one fellow gets a slightly different answer that is wrong, but within 50 percent ...?

MR. MC INTYRE: But this happens on the 6600 under certain circumstances. You have nothing there to help you detect that same kind of failure, either. There is very little parity.

QUESTION: I am sorry, I am not used to a 6600. That happens on a 6600?

MR. MC INTYRE: Sure. It will happen on any machine. Very few machines have that good a parity check, and if you happen to just run into an insidious hardware error, you had better be able to detect it some place during the course of your computation. Otherwise you will have wrong results.

QUESTION: Well, we have some of these energy checks and combination checks, but the IBM machine I am working with has one variant for every eight bits of memory, and I have not seen it fail in that mode. As I say, I was not used to that type of performance.

MR. MC INTYRE: Well, this is not optimum. It would have been nice to have some parity bits in there, but I don't think that kind of problem is of any greater magnitude than the problems people are working with today. As I say, in the 6000 series there is very little parity checking. The only time it is actually done is on the tape where it is checked.

QUESTION: Do you check parity on that disc?

MR. MC INTYRE: The hardware checks it and the programmer does not have access to it. If you hit a fault in the I/O system, the B-6500 which monitors the system, would know it. It would then try to re-read the disc three or four times.

QUESTION: In one sense I am concerned with some of the hardware failures with I/O. When we were running on the 7094 we were pushing around so much data on our calculations that we were getting two bits dropped that would bypass parity at one time. If there was no parity or limited parity, especially when we were moving indexes or any values around, we could get catastrophic answers, answers that were not blatantly wrong, and would get through the energy checks and all of these other things that are in the codes. They would just give the wrong answer by whatever was tolerable within the physics.

MR. MC INTYRE: I guarantee that those things will occur, and that they have occurred in calculations that are published right now. It is a problem. How do you get around it today? I don't know. The only thing you can do is go to the very exotic error-correcting codes, which Burroughs has on some of their equipment.

QUESTION: What I am concerned about is this: if we were going to do some serious ILLIAC forward type calculations with this type of reliability, we would have to program in redundant calculations that will validate our results.

MR. MC INTYRE: Not if you have them in right now in your 6000 codes. You see, the number of computations you can do before failure on the ILLIAC is probably an order of magnitude larger than the number of calculations you can do on the 6000 series without a failure, because the machine does calculations about a hundred times faster.

QUESTION: I am only concerned with detecting the failure in a way that it is not catastrophic. I understand you have a problem, but it is a conventional problem.

MR. MC INTYRE: The ILLIAC has not imposed any ....

QUESTION: I would argue differently. I would say that, in the conventional mode, it is probably not as apparent as it is in a machine that is a thousand times faster with less chance of detecting the error. Then it is up to us to make redundancies in our calculations to be detected. I mean if it does occur, it is catastrophic.

MR. MC INTYRE: Well, I will talk to you about that.

QUESTION: I just wondered on what you based your reliability?

MR. MC INTYRE: These are based on the reliability calculations as prescribed by some Air Force standard document, for Air Force electronic data processing equipment procurement. They are in general quite

conservative, in the sense that I doubt very seriously that the calcu-lations are in error by being too long. They may be too short. We may observe 7 hr instead of 5. Count the number of electronic components, and you know the probabilities--it is a Poisson process, the probability of failure of each component.

QUESTION: If you start failing in 3 hr then what do you do, burn the document?

MR. MC INTYRE: Then we try to recover some costs from Burroughs. In the early stages you may see mean time to failure that is shorter than 5 hr, which probably means the equipment is not shaken down yet.

QUESTION: Okay, but it is not based on some kind of experience data?

MR. MC INTYRE: No.

QUESTION: When will the machine be operational, both hardware and software?

MR. MC INTYRE: If component deliveries stay on schedule, and if de-bugging goes smoothly, it will be early winter. We have scheduled no calamities. It is very hard to turn in a schedule to ARPA which shows a PERT chart including one little block which says "Unscheduled calamity."

QUESTION: What is the past frequency of calamities?

MR. MC INTYRE: We had a horrible calamity with the components of which the machine was to be built, but we lived through it. Originally, the processing elements were going to be very highly integrated and considerably smaller. But a year and a half ago Texas Instruments decided they could not achieve that degree of integration. We had to fall back and implement the processing elements out of circuitry that was not as integrated.

We also ran into a calamity on the memories, but we came out of that very well. We had originally planned to have thin film memo-ries, but they were fairly expensive (17 cents a bit) and were continu-ing to get more expensive. We decided to go to semiconductor memories, and went to two vendors, Fairchild and Motorola. It turned out that Fairchild could build these high performance memories very cheaply, for around 10 cents a bit.

QUESTION: What degree of integration do you have?

MR. MC INTYRE: I am not an electrical engineer so I can't really say. They are dual in-line packages with 16 pins and in general you get maybe two or three gates on a 16-pin package.

QUESTION: I guess that is a little super IC, isn't it?

MR. MC INTYRE: It is very good for this point in time.

QUESTION: When I went to the first ILLIAC thing about a year and a half ago, they were talking pooled MSI.

MR. MC INTYRE: Yes, that was just at the time Texas Instruments said they could not achieve it.

QUESTION: So you have fallen back to ....

MR. MC INTYRE: A more conventional degree of integration.

Mike, why don't you talk a little bit about software?

MR. RANDALL: Well, to start with, we will go from the outside and work in. To remind you of the type of picture that we have now which applies to the software, you might remember the old 1401 had 94 sort of configurations, the idea being to keep the 94 going at full tilt, and do all of your input-output on the 1401. Somehow or other we have managed to get much the same kind of configuration on the ILLIAC IV, except in the middle you have a disc, and you possibly have some equipment hanging out the front. From here on these are completely transparent to the user, so I won't talk about them. It is just as if they were not there. So I am concerned with this, and essentially what happens is you have your line printers and your B-6500 discs, and later on a large memory, and all of the input-output on the B-6500. The B-6500 does all of your addition to binary conversion and back, and the I-4 does all of the heavy number crunching.

If you look at the program, it consists of roughly four parts. It consists first of all of what we call a dot free processor, an ILLIAC IV program itself, and later on a post processor, and the whole thing is tied together by a piece of job control called ILLIAC control language which assures that all of these are done in the correct order.

These are B-6500 programs and this is an ILLIAC IV program of one kind or another. Of course the B-6500 is very good at multiprogramming anyway, and given a whole mix of jobs, some jobs will be having their preprocessing done, some will be having their post processing done. Compilation by itself is done on the B-6500 and transmitted into the I-4, and so on. So what is going on in the B-6500 is a mixed bag of tasks all of the time, preprocessing, post processing, compiling, and most of operating system is in the B-6500.

This control language allows you to open files on the ILLIAC IV disc, to map them in a sort of rudimentary way, although in a reasonable way, to initiate processes on the B-6500 to start programs, to compile, and these programs deliver our result which we can interrogate, and then you know whether you want to go on with the next particular process, and so on. This is a language in which the basic structures that you are dealing with are either files of data or

223

programs. You can ask if this program has been done, if it has gotten through without any errors, so that when you hold it up, it says you can initiate this program, and so on. This is the general outline.

Normally any job will take four programs or four distinct tests to be submitted, and this is more or less divided up.

Now, in the operating system, which is mostly in the B-6500, there is a small section of protective store in the ILLIAC IV. In that protected section is a little bit of the operating system, the only part of the operating system that looks after the loading of programs into the memory and relocates them, and unlists labels and things like that, and looks after the transfers to and from the discs, which it does by transmitting messages through a sneaky little part back to the basic B-6500--to the main operating system--which then does all of the necessary juggling. If a command is wanted here, it will come out of here into the 6500 which then activates the transfer.

QUESTION: Is that spread out over all of the processing elements, that piece of ILLIAC IV memory?

MR. RANDALL: Yes. All programs in the core--if these are your 64-word memories, any program is straight across the memory. There are two instructions per PE memory. Each instruction is only 34, and it just goes along. It must start on the late word boundary, or something, out there at the beginning.

Thus the input-output is rather a complicated business going through here, and as Dr. McIntyre pointed out, the ILLIAC is good for programs that you can load up the disc and then work on the problem on that disc. A way of looking at ILLIAC IV is that it is essentially a machine that transforms the contents of the disc in some way, and this is the way it should be used.

The operation that is in the B-6500 is essentially a set of subprograms that talk to the barriers in the CP, which is the barriers operating system. In structure you have the MCP (the master control program), which is sort of the king of the castle, and then you have all of your particular data sets and other processes beneath that.

One of the essential consequences of this, or this kind of approach, is that the ILLIAC IV turns out to be the master, and the whole configuration is purely to keep ILLIAC IV going at top speed. It is demanding all of the time, demanding jobs all of the time, and these are being fed to it by the B-6500. That is why I say it is an exact analogy to the 1401.

Are there any questions on the operating system?

QUESTION: How complicated is the operating system language likely to be? Presumably you are going to have to supply some of this to run your program.

224

MR. RANDALL: It is not very difficult. I think the most difficult parts are those which specify the file, provided you can map on the disc, and you can map areas, or rather either segments or records of a file, in milliseconds around the disc. You prepare a program. If you clear a program X which delivers a parameter I when it is finished, then later on you come back and do this kind of thing. If you clear program X then you can execute it, and if it delivers a reasonable result, then you can go on with this. That is more or less what it is like. I don't think it is any more complicated than the 360.

QUESTION: The 360 does just that, goes through job steps, and you have the options.

MR. RANDALL: Yes, but that is not all of the mumbo jumbo of flow charts and everything behind what is here.

QUESTION: I suppose the main concern would be will it undo any of the previous languages that one might have learned, or will it be in conflict with any of the other languages?

MR. RANDALL: You mean job control languages, or which languages are you talking about?

QUESTION: Like FORTRAN. If somebody puts a statement of that sort into the machine, will it be likely to be totally rejected, or will it be incorporated into an overall language, a super-language, so to speak?

MR. RANDALL: No, we are not incorporating into super-language. Are there any more questions on operating systems? If not I will go on to languages, and answer your question in more detail.

QUESTION: I have a telecommunications question. It seems to me that would be sort of a third area. I am concerned now with the ARPA net, and I am the user in California. My first question is what software is being developed for that telecommunications problem? That would be B-6500 software, and the net, and the whole concept of telecommunications.

MR. RANDALL: We have a software group that is working on communications to give you exactly the same kind of access you would have if you were on the site, really, if you sit down at the machine and type your demands in.

QUESTION: So you have a group working on that.

MR. RANDALL: Yes. They have just started. As soon as we know more about the kind of equipment, we will talk about it. It is much the same as using any other machine.

QUESTION: There are some problems with hydrodynamics codes, like the output. They are basic problems that one has to address himself to. Maybe you will have a huge data pile and you just send graphs back.

MR. MC INTYRE: Once again, the fact that ILLIAC IV is there, you see, does not complicate the problem. It is the same problem of trying to use the B-6500, a very conventional machine, remotely for the ARPA data. Right now 1108 and 360 machines are being used remotely out on the West Coast through the ARPA net. All we have to do is write the software so you can use the B-6500 remotely, and then you can use the ILLIAC IV.

QUESTION: Is that use of the 1108 and the 360 in the ARPA net operational now?

MR. MC INTYRE: Yes, the 1108 and the 360 do work.

QUESTION: So you just become a subset of that telecommunications system.

MR. MC INTYRE: Right.

QUESTION: Does the remote thing include fast output via microfilm or something like that?

MR. RANDALL: Yes, there is microfilm being provided, but not remotely.

QUESTION: Well, I don't mean remotely, because I can think of printing out tens of thousands of pages on this thing.

MR. MC INTYRE: There will be a microfilming device on the ILLIAC system, but there is no way we can send that at those bandwidths back through the net. The net only has 50,000 bits per second transfer rate. You may have to settle for mail or courier.

QUESTION: You would be very selective about your printout. What is the printer capacity?

MR. MC INTYRE: The 6500 will have 2,000-line per min printers on it, but the microfilming device is about ten to twenty times faster than that.

QUESTION: Just two printers?

MR. MC INTYRE: Just two. Printers are very expensive.

MR. RANDALL: It is more reasonable not to get 64 tons of line printer out, but to look at it more selectively and just destroy the files once you have the results you want.

QUESTION: How big a mass storage will there be on the B-6500? You said a trillion bits.

MR. RANDALL: Yes, 10^{12} bits, but that will be 18 months or 2 years away.

QUESTION: I was thinking of a system where you keep data tapes at the B-6500 place, and just spit out at the graphical terminal or something.

MR. MC INTYRE: At the graphical terminal, if it requires less than 50,000 bits ....

QUESTION: It depends. Is that just one 50,000-bit channel?

MR. MC INTYRE: Yes, there is some talk about duplexing it, but initially there will be just one.

QUESTION: So two people won't be able to get graphical output at the same time.

MR. MC INTYRE: That is right.

QUESTION: I would say mail digital tape or computer tapes or something, to start off with, because you can't do a 3-D problem on any other local machines. You can if you want to wait two weeks for the answer, and then you have this machine time problem.

MR. RANDALL: I think there is a bit of adjustment in using the ILLIAC IV because of the large amounts of data that are likely to get spewed out at you. In the present operating system, there is a tendency to run one job to completion to minimize the number of transfers. The idea is to put the job in and wait for it to finish before you bring another one in. This is a pretty old fashioned kind of operating system.

QUESTION: I guess to summarize my question then, that telecommunications thing is operational now, and the work is being done to solidify the ARPA net?

MR. RANDALL: Yes, that is right.

QUESTION: As I say, most users will be remote, and no matter how powerful the ILLIAC IV is, if we can't get into it, it won't be very helpful.

MR. MC INTYRE: Right now people in California are using an 1108 in Utah, and vice versa, people in Utah are using 360's in California.

MR. RANDALL: At present there are only two languages available that are being used in conjunction with the simulator. These are the language called ASK, which is the equivalent of the machine language, and the language called GLYPNIR, which is a little narrower, but is more user oriented than ASK. ASK is just like any simple language; it is a list of instructions and addresses with labels in front of them. It has some pretty powerful macro possibilities. These are sort of the replacement or defined statements. You can really use the machine efficiently with this.

GLYPNIR, on the other hand, is more like ALGOL, where you can use some varying statements easily. You can use a vector 64 lists long. This is one of the problems we are finding in the software. There are two basic problems. One of them is finite with the machine, and if you want to do something really funny with an array that is 65 long, then you have to find some way to reconfigure it to use the machine properly. Since you have to reconfigure your problem anyway just to put it into sort of parallel algorithm, this is probably not much of a restriction. On the other hand, in high label languages, this difficulty of doing the reading in a reasonable way exists, but once again there is the same problem. If you want to do a lot of routing, you have to say the boundary conditions if you are doing some equation problems. So you can easily switch your boundary conditions on and off.

We are talking about a FORTRAN, and I think Barrett has to do it, but we have not done anything, or we have not had any definite answer from them, but we would certainly like to be using FORTRAN.

What does parallelism mean? What does it look like in a highly different language? The idea is that you refer to arrays in the ordinary way, for instance, a two-dimensional array used as we were talking about before, again we refer to the whole row. I won't finish this, and you all probably know it, but what does the asterisk refer to? This is a parallel expression. It is done right across the array U at one fell swoop, if the array is not more than 64 long. What does the asterisk mean? The asterisk refers to those PE's that are turned on, so essentially in front you have to do this, which equals some particular vector, and this can be written thus, and this actually turns on in this particular case all the PE's except the two outermost ones, just for example. Then this instruction will be done in one bank, and except that U-0 and U-64 will be turned off, and they won't figure in the calculation at all, that is, even when you pick them up.

QUESTION: Do you have an indicator for the identifier I?

MR. RANDALL: Yes, you have that on the outside.

Talking about how to design FORTRAN, there is this kind of thing: Now, suppose this I was more than 64 long, then we have two choices before us. You either do this expression twice, one for the first group of 64, and then for those in the next group of 64 across, and then for those left over, in which case you would also be fiddling around in the background in showing that you do do it, or else as we have chosen, it is up to the person himself to explicitly structure stuff so it is only 64 wide. Of course, the other solution is to just do one strip of 64 down the code, and then the next strip of 64 down the code, and then the next strip, and so on.

If your code is cross referenced so that something on the left also appears on the right, then with the strip thing you are not going to get the right answer anyway, and you ask yourself what does this

statement mean if you are only doing it with strips of 64 down rather than strips of 64 across. What is the subroutine intervening if you are doing it in strips of 64 down rather than strips of 64 across?

This is why we have chosen to limit the software to a width of 64.

QUESTION: I am sorry, you said I. You meant the star. The I can go from anything as long as the star goes from one of these 64 strips. It does not matter. That is statement No. 6.

MR. RANDALL: And in statement No. 6 I am talking about the width.

The same with the routing, if you look at the store this would be row 1. You want to load this row and move the long one, then you should have moved this down here. Then you load this down here somewhere else. This is why it is much easier to limit our row to 64, and all of our structures are made in multiples of 64. To expect the program to do it at this stage, because of these machines ....

MR. MC INTYRE: To clarify what he is saying, you are not restricted to work on vectors that have only 64 components. You must program it in groups of 64. You must actually write hexcubital statements in groups of 64 or less, so if you are working on a vector that is 100 long, you write a program addressing the first 64 components, and then the next 36.

MR. RANDALL: Are there any more questions?

QUESTION: In that example up there, you are just saying to allow I to range over anything you want. First you only have 2,048 locations in any memory, don't you?

MR. RANDALL: Yes, I did mention the 2,000. That is about the limit.

QUESTION: Here is something that again you have to remember the finiteness.

MR. RANDALL: Yes, and you nave to keep your rates filled.

QUESTION: What kind of I/O is in GLYPNIR? My same question about buffer in, buffer out.

MR. RANDALL: We are going to write the I/O for GLYPNIR this summer, so it is not entirely finalized, but the kind of thing is that you will read from nine files with reinstatements into particular arrays, and there will probably be a big structure behind, so that you can build up a queue of I/O requests, and then interrogate the big structure behind to actually find out whether that has been completed.

QUESTION: So it will in fact look like the buffer in, buffer out. In other words, there will be some kind of conditional telling you that the I/O has been completed?

MR. RANDALL: Yes, that is right. It will either be a condition telling you it is completed, or there will be an automatic holdup if it has not, one or the other.

QUESTION: That holdup will have to be on a conditional statement, because the machine won't know it is not completed if it goes to that array, unless you flag it, which will make the execution time slower, because every time it tries to get to the array, it has to check to see if it is flagged.

MR. RANDALL: That is right. In the software, this is sort of the region of greatest compromise, because you are in a position where you can rely on the individual program an awful lot; if you don't, all hell can be let loose.

QUESTION: This is a tremendous problem with IBM. They refuse to take that CEC approach of buffer in, buffer out, and for the hydrodynamics codes, that is what you have to have. They still don't have it on 360-OS in the FORTRAN mode, and every time we switch machines we have to write assembly language version of this buffer in, buffer out.

MR. MC INTYRE: I do want to mention some benchmarks that we have done to give you some idea. The Weapons Lab commissioned a couple of benchmark studies on two different codes. One was called HEMP and was a high altitude EMP calculation. We found for the central compute portion of the code, the ILLIAC IV was 40 times faster than the 6600. The other code was SC, which was a two-dimensional version of an EMP code, and that one turned out to be 90 times faster than the 6600. Yesterday you saw me give those figures which say that in the equation of state calculation, in this funny comparison, ILLIAC IV could do 64 equations of state calculations in the time required for one long equation of state calculation for the 6600. We did do a benchmark for the National Security Agency which found that the ILLIAC IV was about 80 times faster than the Burroughs 8500, and 265 times faster than the 360-65 on the same problem. Those were all exclusive of input-output. Those were the central computation portions of the computations.

QUESTION: Would you explain how you do these benchmarks? Do you simulate the ILLIAC IV part of it in some sense?

MR. MC INTYRE: Yes. These were all simulated and all written in assembly language. Once again we counted the clocks, which is a rather conservative estimate of the time required for ILLIAC IV because it does not take into account the overlap between the control units or the overlap within the process.

COL. RUSSELL: For the next part of our program, I have asked Bill Whitaker from the Air Force Weapons Lab to give us a very short discussion on reconfiguration problems.

MR. WHITAKER: I represent in some sense the opposition, that is, I don't build machines, and I don't, for this purpose anyway, write software. I am a user. I write big hydro codes; I run quantum mechanical calculations. I run as big hydro codes and as long quantum mechanical calculations as anybody else in the world. Therefore, I have a special interest in the ILLIAC IV. The interest is to get numbers, just numbers. That is all I care about.

I would like to talk very briefly about the ILLIAC IV in that sense. ILLIAC was originally started as a great advance in the state of the art, a big step forward, a machine that was hundreds of times faster than anything we had at the time. It was started by ARPA as an advanced development, and was in some sense unique, that is, ARPA was sticking its neck out and making a big step forward where it did not look like anybody else was really going to go this far.

Other commercial organizations have in fact followed along very closely, perhaps to a certain extent prompted by the success and interest in the ILLIAC IV. In any case, the ILLIAC IV is not going to be as unique as it might have been, but I suppose you could not really have expected that. There are going to be other machines on the same time period that are of comparable speed: the Star, the SPS, and Texas Instruments has a machine they are proposing now. But the group here, I assume, is specifically interested in the ILLIAC because it will more preferentially have access to the ILLIAC. These other machines are not going to be all that available, so perhaps we still have a unique feature here.

There are going to be a lot of problems for someone who needs numbers, in spite of what these other gentlemen have said. They represent another area, and they are working hard to get their system working. They are working hard to deliver something that they can be happy with, but I must allow again that I am the opposition, that I am not necessarily going to be happy with what they are happy with. My apologies.

First off, let me tell you a war story, let me tell you a couple of war stories. The last big change in machines was going to the 6600. Now, the 6600's were not a terribly impressive change in the general structure of machines. They were impressively faster, but still to a certain extent, as far as the outside programmer was concerned, a serial machine, so there is nothing particularly strange about that. You still used your regular FORTRAN programs, and in fact, a FORTRAN program from your 1604 or your 7044 or your 7094 could, if it were well written, run immediately on the 6600 in principle. It didn't always work that way.

231

We got Serial 6 of the 6600's, the sixth one that was made, and we had the opportunity of checking out the machine at the factory as it was being assembled. We found that there were some difficulties. The machine, as far as the outside programmer was concerned, was serial. In fact, internally to the machine, the arithmetic units were parallel. This was something new and different as far as machines are concerned, that is, it didn't take one number, add another number to it, think about it for a while, and then multiply it by something else. It was liable to do it all at the same time, and timing became extremely important in the 6600 in a sense that it had never been important before. As a result, and my apologies to CDC if they are not here, it was possible to run the normal checkout programs, going through memory and taking every number and adding it to something else, or reversing the bits or something like that and it worked fine. Obviously it worked fine. That is what the CE's were sitting there adjusting it for.

However, when we loaded on a program, the simplest of all possible programs, a 1-D Lagrangian hydro code, which was our basic test problem, it would not work. It would not even compile. The whole thing did not work, and it took us a long time going through the machine and through the software to get even the simplest program to work effectively.

Now, this was in spite of the fact that all the checkout programs, all of the normal sorts of things that the engineers do, worked perfectly. The problem, of course, was timing. It not only depended now upon every arithmetic unit being able to work, that is, the add unit adding and the multiply unit multiplying, but the results were very sensitive to timing, very sensitive to the order of the instructions, sensitive in some cases to the bit patterns that you fed into the units. This machine was more complicated than we had ever looked at before.

Well, we were very proud of our machine--our machine we checked out at the factory. After it was delivered, within something like ten days, it was operating 24 hr a day seven days a week, and giving right answers.

There were other machines, for instance, Serial 2, which never worked. They finally sent it back to the factory and redid it. Serial 4 I believe was in Geneva at CERN, and they paid penalties on that machine for 6 months, which was the contractual limit, because they were unable to run a program on it.

Now, unfortunately, I propose that this ILLIAC machine we are facing today is more complicated, is more difficult, and it certainly is not going to be easier to get to work. The problem here again is the situation where the people who designed the machine and who are working with it are going to make it work for their sorts of things. The people with big production programs simply have different problems. They are going to run into situations that the machine designer has never thought about. They may be pathological situations. They may be things that we

should have avoided. Nevertheless we are going to have them. We are going to have strange things in the software, strange things that we have managed to live with for years, and all of a sudden they don't work any more. There are going to be a number of such things. It is going to take a long time to shake out the machine to a state satisfactory to those who want numbers.

Let me briefly review what I consider a production program. We have timing estimates here for the very, very inner loop. These represent, I think, in no case perhaps more than a thousand instructions or a couple of thousand instructions, FORTRAN instructions. Big production programs are in many cases a lot bigger than that. There is a lot of input-output. There is a lot of handling of data in strange ways. There are graphics. There is setup. There are lots of things that are not proper to go through in the timing operation for this, but that we are going to have to go through eventually. In fact, if your inner loop works, that does not really help. You still have to have answers out the far end, and that is going to take a while.

We are talking about big programs. There is no point in using the ILLIAC for something that runs a minute on a 6600. In fact, I suppose there is no point in using ILLIAC for something that runs an hour on the 6600. We are talking about programs that run tens of hours, hundreds of hours. Not all of you may represent programs which run hundreds of hours. Perhaps some of you wish to represent programs which run hundreds of hours, but have never had the opportunity. But we certainly are talking about such big programs, and programs that will be expanded to the equivalent of such. Otherwise, there is no point in trying to use a new machine. In fact, there is going to be a fair amount of overhead, a fair amount of effort, even under the best of circumstances, to getting on an ILLIAC or on to any other large new machine. To make this effort worthwhile it obviously has to go to a point where you are talking about running an amount of time on the 6600 or say an 1108, equivalent to the amount of money you are going to spend in getting on the new machine, or perhaps twice that much, since it is not only just the effort you spend, but it is a lot of trouble, too. Perhaps you are talking about spending tens of thousands of dollars adapting these programs. That becomes hundreds or thousands of hours of machine time for 6600's or 1108's. Thus you are talking about the very big programs, programs that either are entirely consuming machines now, or programs which you would like to write to entirely consume machines, programs like SHELL or CORONET, NIXON, quantum mechanical codes, etc.

In that light, certainly the machine will be more useful if we are able to greatly reduce the amount of effort necessary to put already working codes on this new machine.

Let me take a census, if I may, among the group here. Of those of you who have active interest in a code of one sort or another, how many of these codes are written in FORTRAN? Okay. How many are

written in ALGOL? Some other language? The prosecution rests. FORTRAN
is the language of large codes. That is true among this group, among
the other DOD laboratories, and now it is for the most part true among
the AEC laboratories also.

We are therefore, I believe, constrained eventually to have an
operating language that looks like FORTRAN. Sure, any of us could learn
any other language we want to. There is no question about that. We are
all smart fellows. There are only a few hundred instructions anyway.
We ought to be able to do that in a few months. Nevertheless, it is
bother, and if nothing else, it is a lot of punching of cards. That
bother is going to considerably increase the effort and the expense of
putting an already working code on any machine. It is therefore going
to reduce the number of codes put on the machine, and that is certainly
not what we want.

Is the ILLIAC unique in this? Certainly not. The other new
types of machines have very much the same sorts of problems. In other
words, for instance, the Star, if I may mention the name in the absence
of the CDC people, is a pipeline machine. You want to line up all of
your arrays as vectors and feed them through the pipeline. Well, you
don't have an instruction for that in present standard FORTRAN. We are
going to have to make an instruction. We are going to have to extend
the language. For the same reason we don't have an instruction in
FORTRAN that says line up all of these processors and do them all simul-
taneously. Actually the problems between those two types of machines
are very similar. The sorts of things that are going to have to be done
are fairly similar, and except for perhaps the uniqueness of the 64 in
this machine, one can imagine very much the same sort of language is
being used, which is not to say Livermore is not having their own
problems with languages. They certainly are.

Nevertheless, to get a program on the machine, ideally what
you would like under present circumstances is to take an operating code,
and it is extremely important, at least for the next year or so, that
it be an operating code. That is, you would not like to start off from
scratch writing a code and assemble it first on the ILLIAC IV. If you
don't know the code is working, then you don't know if the machine is
working, and that is a very vital part. We have to first go with codes
that are thoroughly understood, working well, and for which we have run
exactly the same problems on another machine, and have exactly the
answers. Otherwise we will not be able to determine difficulties in
the coding, in the assembly, and in the hardware.

We would like to take a code which is written in FORTRAN as
it now stands, make such minimal modifications as are necessary in both
the logic, which in many cases may be necessary, but not all, and in the
extension of the language to bring in the unique features of the ILLIAC
or whatever machine you may have, and then run that directly.

234

Well, I am now not so much the opposition. In fact, the University of Illinois has a proposal which sounds very much like that, but it does not stand as a system at this point. Nevertheless, we recognize that that is the ideal sort of thing. But it is going to take a while, and what can we do in the interim?

I think you will all agree that it is important that we don't send off a bunch of programmers or physicists, as the case may be, to the ILLIAC, each by himself to learn the new and exciting system, to make all of the mistakes he can (which is how you get to be an expert, by having made all of the mistakes, each individually). This only compounds our number of mistakes. In unity there is strength. Somehow, the users, the people who want numbers, are going to have to stick together in some way, organize or be organized perhaps. There will have to be a firm program by the users to get programs on the machine, to check out the software and the hardware, to learn and document the mistakes, to bring people along as they come into the system, and most important, to design from a user's point of view what you would like to have in terms of an operating system.

As has been mentioned, the people who design machines or the people who market machines are not always those best fitted for designing specifications of such an operating system. We have lots of examples of it. Everyone here I am sure can cite several things he would like to do with his machine or his software, things that are important, but a single user does not really have much chance of getting such things initiated, particularly if you are fighting all of IBM or CDC or Burroughs or GE or any of the others. It is a big system, and you have trouble, and you are not going to make it. You are not going to write to Mr. Watson and say, "Gee, I need this buffer in and buffer out." It is not going to do you any good.

UNIDENTIFIED: We just wrote him. That was no problem.

MR. WHITAKER: That is right, but you don't have it. We are very fortunate at the Weapons Lab, for instance, to be able to govern our operating system in some detail. We are now left out in the cold because we have a unique operating system and all of the rest of the world operates on something else, which is some difficulty. Nevertheless, ours is better, but there comes a point where even being better is not good enough. We were very successful in operating for five years under those circumstances. The reason was that we had a small number of people, users, mind you, who completely control the machine. We were together, unity.

That is no particular reflection upon the people who are trying to do the software. They were trying to do the job as they thought best, and they were doing it. It is just that their priorities were not the same as ours. This is perfectly reasonable.

But the purpose of the machine, the purpose of the ILLIAC, the purpose of our 6600's, the purposes of all of your machines, as far as you are concerned, is to get numbers. As long as we can stick together and as long as we have direction in that sense we should be able to make this a good operating machine. But it will take a while. Thank you.

QUESTION: Before you leave the stage, what would you suggest for sticking together? Everything you say is true, and we all know it is true. What are you suggesting?

MR. WHITAKER: I believe that one organization should be set up which would have control, if you wish, in this sense, of the users, not a committee, not a cooperative group or anything, but a director.

MR. MC INTYRE: You mean one that is not set up now.

QUESTION: But how does he suggest setting it up?

MR. MC INTYRE: The machine belongs to ARPA, and the Department of Defense can designate.

MR. WHITAKER: That is right, the machine belongs to ARPA. It is the responsibility of the Department of Defense, and the Department should exercise that responsibility.

MR. CHERRY: Who has been making these decisions about what happens to that machine, or is it just sort of haphazard? I am sure it is not.

MR. MC INTYRE: It is administered the same way any ARPA contract is. You have to understand that this contract is progressing from a research phase into a service phase, so just as administration of a research contract is a little different from administration of a purely service contract, what Bill is saying is that you have to administer research one way with a little bit of freedom for the researchee, or researcher, and you have to administer service another way. You just change the orientation.

MR. CHERRY: But I think you are saying that you are going from a research effort to a service effort.

MR. MC INTYRE: Yes. Of course, as Bill points out, those first few months are going to be--you know, we would be naive to say we are going to stop research here and run 24 hr a day. There will be a transition period.

MR. CHERRY: One of the things that bothers me as far as the codes I would be running if it becomes our concern, after a while they have a tendency not to run efficiently on the machine. You know, they run for a few hours and then it is just not profitable to run them with that configuration very much longer. You have to take the problem off and do things to it to get the problem back to the condition where it will run.

MR. MC INTYRE: Human interaction is required every so many time cycles.

MR. CHERRY: You made the statement that you would like to run 40 hr.
Well, I run problems 40 hr, but I don't run them 40 hr at a time. I
have to look at the thing at certain points during the running of the
problem. If there is a week's delay between the time I run the problem
on the ILLIAC and the time I see the results, it is going to take me an
awfully long time to do the problem.

MR. MC INTYRE: I would surely hope there would not be a week's delay.

MR. CHERRY: How much time delay is there going to be for a CD output?

MR. MC INTYRE: If you are working through the ARPA net, you can get
back considerable information through the net. It takes something like
20 min to write a half inch tape completely full of information, which
you can then process any way you would like to at your installation. So
in 20 min you can get a whole tape's worth of information, which is con-
siderable information. If you want plots, it may take the time required
to develop the film and put it in the mail and get it to you. In that
sense it may be two days. But you certainly could get snapshot dumps or
partial pictures much quicker than that, like just a few minutes.

QUESTION: I get the impression from what I heard yesterday and today
that this higher level language is based upon ALGOL.

MR. WHITAKER: It is. As it now stands, it is based upon ALGOL.

QUESTION: Now, here is a decision that has already been made.

MR. MC INTYRE: That is right, and that is a reflection of the research
orientation of the project.

MR. WHITAKER: It is, however, not a sufficiently higher order language.
It is not going to do all of the things for you, even if your program is
written in ALGOL, but you are not stuck with it. What I am pointing out
is that this is not the final thing but is a reflection of the earlier
portion of it.

QUESTION: I take it from what you say you are years away from it.

MR. RANDALL: I endorse everything Bill says about forming a users group.
How many people here have been approached about the FORTRAN design?
There are many more, and they have great interest in what we are doing,
but we get no input from them whatsoever, so we just have to go and make
arbitrary decisions. I agree with Bill that we must have a users group
if we are going to even give you 10 percent of what you want, but you
are only getting about 1 percent now.

MR. WHITAKER: Yes, but we have had a certain amount of that interaction
from the users group, as you say, before.   However, it has been a
group without responsibility.

MR. RANDALL: Yes.

MR. MC INTYRE: Without authority.

MR. WHITAKER: And that has to change.

QUESTION: I think we have a problem in de-escalation, though, because originally the design specifications were going to satisfy the users with a high-level FORTRAN language that would look ahead, and in fact initiate parallel operations in the standard FORTRAN. I think the people at Illinois discovered they could not do that, and to allow the machine to be in use the first year, they provided this GLYPNIR thing. If we want to get moving as fast as we can on 3-D codes, I think the responsibility is with us, the users, to conform to what is available now for the first year and a half, and at the same time to form some authoritative committee that will hopefully allow us to have what we want, say, two to three years from now, which is behind the original ARPA schedule, but that is not what we have now.

MR. MC INTYRE: Now, wait a minute. You can't form an authoritative committee. The machine belongs to ARPA, and Information Processing Technology in ARPA. Before you can think about one-sided action, I would suggest that those people be considered. They are very reasonable people.

QUESTION: Actually the first part of my statement is that if we want to do 3-D codes and not run 100 to 300 hr on existing machines, we have to do it this way. We don't have any choice, because the feeling I get after watching this thing over the past two to three years is that the parallel nature of the problem is not tangible to creating a universal language, and we have to have GLYPNIR or we won't be able to allow ....

MR. MC INTYRE: GLYPNIR is not an end in itself, and nobody has designed it to be. It is to hold us over. Remember what GLYPNIR stands for. In old mythology there was a wolf which ran around eating up people and doing bad things, and they could not hold him. They made a chain out of very exotic elements which held the wolf for a short period of time. That chain was called Glypnir. You can draw your own conclusions.

MR. RANDALL: I wonder if I could make one or two remarks about language development, and the kind of philosophical problems that we are up against. Ten years ago when they started producing languages like FORTRAN, they did so basing them upon a well founded symbolic technique used for ordinary calculations that was already there for them. All they had to do was to take the language of mathematics, put in a few incomprehensible punctuation marks to confuse everybody, and present it as a programming language.

  Now, when you come to pipe lining or the ILLIAC IV kind of parallel processing, there is no explicit formulation, no explicit formalism on which we can rely. So just putting in a few

238

incomprehensible punctuation marks is pretty difficult. We have to go through about ten years evolution in computing languages in about 18 months or two years, and we have to try many ways. GLYPNIR is one of those ways. This is one of the reasons why we are referring to it that way, apart from being the chain that keeps the wolf at bay. It is also a pretty valid experiment in parallel languages.

If I could refine something that I think Bill said, the idea of taking a serial code and making minimal alterations so that it will run efficiently on ILLIAC, I agree that one should make minimal alterations, but I think the changes in the algorithm make it a pretty big minimal.

MR. MC INTYRE: Well, hopefully not the entire code.

MR. RANDALL: No, no, just the central loop.

MR. WHITAKER: I think actually what you are talking about in those cases is making alterations in a portion of the code which may represent 90 or 95 percent of the actual computation, but not 95 percent of the cards.

MR. RANDALL: No.

QUESTION: Just to fix the idea, would you mind telling us what TRANQUIL was originally supposed to do?

MR. MC INTYRE: TRANQUIL was a language, which was structured after ALGOL, incidentally, that was designed to hide the architecture of the machine from the user. In other words, the user did not have to know that he had 64 processing elements. He could assume he had any number, and the user did not even have to worry about the amount of core he had, because input-output would be taken care of for him. We got to the point where we could start to test some of the object code, and it turned out that the code was running from a factor of ten to a factor of twenty slower than similar codes which were in assembly language. What we were achieving was to get a machine that was 100 times faster, and then immediately giving back a factor of 20 through the software, which says we have a factor of 5, which we regarded as unacceptable.

QUESTION: Is it too much of an oversimplification to say that the symbolism of arithmetic and algebra is to FORTRAN as the symbolism of matrix algebra was to TRANQUIL?

MR. MC INTYRE: That is pretty close to it.

QUESTION: I dispute your point about the symbolism is not there to work on. It may not be easy, but you have matrix algebra.

MR. RANDALL: It is there, but it is not cut up. You see, you talk about multiplying the matrix A times B. It would be possible to produce a machine with the architecture that would do that, but the problem

239

with TRANQUIL was that it was not applicable with the architecture, and tried to disguise it. In our software thinking, I think we are being sort of reactionary at the moment. We are going the other way in the FORTRAN that we are designing. If you know about the architecture of the machine, you can write pretty efficient FORTRAN. The idea is that you want the AB, the multiplication of two matrices, but computationally you are actually working on rows of elements and rows of columns, and to try and express these in the simple ... if you have a machine that needs what the matrix might program completely, and you put that foremost, that falls in between these two stools.

QUESTION: Bill asked earlier how many people's codes were programmed in FORTRAN. If he had asked that same question in 1961 or 1962, there would have been a lot of people programming in machine language, because in those days you were paying a lot for the FORTRAN. I think that is the situation you are in now. You don't start out with an efficient, high-level language, but you have to start. You have to get a lot of people working at it. As Bill says, you have to learn by making mistakes.

MR. RANDALL: Believe me, as far as we are concerned, the object of TRANQUIL is not abandoned. It is not dead. Our present thinking, and it might change, is that if we were going to have a TRANQUIL language, then both TRANQUIL and the machine architecture have to change. This is one of the attractive things that should be done, that we try and marry the two, and then have to put out new proposals for a new machine.

QUESTION: I think Bill's point, though, about taking a code that is 8,000 statements long (of which we have three of four sitting around) and just putting it on the machine is very well founded. One will have to pay particular attention to the details of that code. On the other hand, if you look at that code and calculate which FORTRAN statements are executed 99 percent of the time, you will find 1,000 FORTRAN statements that are using CPU time 99 percent of the time. So you can't be naive when you go into ILLIAC IV. What I think you are saying is that we really have to take close looks at it, and there are no panaceas. This is the message I am getting, which should be given, that we have to look at those programs and decide which part of the program can be converted and which part can't, yet be able to run the whole program within two years from today, or whatever the time scale is, or six months. I am talking about the fact that hardware may be ready, but we may not be able to get the codes converted properly because of this problem that has been pointed out, and that is part of the issue. Of course, with the present level of software, that inner part will be in TRANQUIL, and then of course you have boundary conditions. Your little example pointed that out yesterday with the two boundary conditions, you see. One little statement in a very simple one-dimensional diffusion equation turns out to be 90 percent of the code, plus initial conditions, in a large scale hydro code or quantum mechanics code. This is the problem we are faced with, and I feel that a part of the responsibility for it still lies with the user himself to understand what he has, as well as

with the ILLIAC people. I don't know how to get the two married, except that I have tried to point out that you need more authoritative mechanisms.

QUESTION: In the evolution of ILLIAC were you considering the language and software at the same time as the hardware, and developing them in parallel, or did the hardware get the head start?

MR. MC INTYRE: The hardware had a modest head start, but it was the chicken that came first. It soon laid the egg. The software went along with it.

QUESTION: There were very early papers on the software. They were considered side by side here four years ago.

MR. WHITAKER: The trouble there of course is that the driving force on the hardware and the software must be different, and it has been.

QUESTION: Does the Illinois group plan to convert any representative code that you know how to do better than anybody else? Say in ASK, which would be the most efficient.

MR. MC INTYRE: We could, probably either in ASK or GLYPNIR. It turns out GLYPNIR is quite efficient. We could if the manpower were underwritten in some way.

QUESTION: That is not a current problem.

MR. MC INTYRE: No. There are proposals to ARPA to underwrite the manpower required to do that kind of thing, and we continue to talk about how much is an appropriate limit.

QUESTION: I would like to review this other point about telecommunications, because I look at a system as being like a series of windows. One window is slightly opaque, and you can't see through no matter how clear other windows are in the system. As I look at the ILLIAC IV, I feel that the weakest link, or the weaker part, would be this net or this telecommunications aspect when used with large-scale computational physics codes, particularly where the output is a problem or where you decide to get graphical microfilm. In some cases, I guess if you have one code feed another, you may have an input problem. You mentioned 20 min for tape, and I am concerned with the efficiency of using a whole system like this. I don't feel it is being done now. The telecommunications terminals are used now for small engineering calculations.

MR. MC INTYRE: It is new, and using big machines remotely is brand new. It has just been going on for the last few months, and clearly it is not as good as if you were at the site. If you want to come to the site, that is fine. There will be offices available for you, and you can come. Presumably during a heavy debug period, you would probably want to be there. However, after the code is set up and is debugged, and all you

have to do is change input parameters, or make minor modifications to the code, I don't see why that can't be done through the net, either mailing the output back to you in two days--in some installations you know you only get one run a day anyway--or sending snapshot output back to you through the net.

QUESTION: Ten years from now the data files will be so sophisticated, and people will all have graphics terminals, and I can see it. For the next couple of years, however, I do see a problem with output. I don't want that portion of this entire system to be overlooked.

MR. RANDALL: For output, I don't know whether you have considered this, sending the bit pattern down and being transformed and getting it back transformed, and then doing all of their normal work on it?

QUESTION: Yes, that is a thought, at 20 min of real tape. Typically we run 100 hr on our computer and create maybe 50 reels of tape, because the output is in tape form and it is never even put out in human form until the guy looks it over and says, "Well, I want to see Zone 35 of the histories."

MR. MC INTYRE: Look, what do you do with that tape? You can't read it off that tape. Actually that tape is just some sort of archive or reservoir store. We intend to have in that trillion bit store at the machine this archive storage. You process that tape at fairly low speeds on a conventional machine to produce some sort of output. There is no reason why you can't process it using a conventional machine, either the 6500 or your machine to the ARPA, working out ....

QUESTION: Yes, I said 100 reels of tape at 20 min a reel, which I guess is okay.

MR. MC INTYRE: Process it at the site, perhaps even using the ILLIAC. Constructing contour plots there is no reason why that can't be done in parallel. As a matter of fact, we have done it.

QUESTION: That is something to keep in mind. Will there be provisions for the software to dump part of that disc on tape?

MR. MC INTYRE: Yes, ARPA has already given a contract to a fellow to handle that.

COL. RUSSELL: I think this has been very useful this afternoon. I would like to talk for just a moment about the Nuclear Monitoring Research Office responsibilities as far as the ILLIAC goes in ARPA. The office owns roughly 5 percent of the machine time that is going to be available when the machine comes on line. We have an interest in reconfiguring codes to use that particular time to meet some of the various research objectives we have in our office.

The general outline of the program we are going to use follows essentially what Bill was talking about. We plan to have a small group to provide technical monitorship of the reconfiguration of the codes. This group would consist of a team of people who are familiar with the machine, familiar with the programming problems, and familiar with the mistakes that have been made. This group would be the technical point of contact, if you will, for the contractors involved in our particular phase of the reconfiguration work.

We hope to start this particular effort in FY 1971. In addition, we hope also to encourage several of our military agents to work with their codes. We will give them a certain amount of machine time, which as Dave has pointed out is worth about $1500 an hour, to attack their problems. We would hope that they could make a parallel effort with us in reconfiguring codes of interest.

For those of you who have codes that you think are particularly adapted to this type of work, I would be interested in your comments, a note or a general outline of why the code is suitable for reconfiguration to the ILLIAC, the length of time it takes to run it, the type of problem it can attack, and why in effect it is in the best interests of the United States Government to reconfigure this code considering the cost that is going to be involved doing it.

That essentially is the program we have. Are there any comments or questions?

QUESTION: Are there going to be any users' guides available?

COL. RUSSELL: The University of Illinois has a large series of publications out now. Are you familiar with them? If you talk to Dave, he can tell you how to get hold of this index of publications that covers a multitude of subjects concerned with the machine, that is, both programming and I understand the hardware, too, is that correct?

MR. MC INTYRE: There are manuals, and we would be happy to give you some. In addition to that, there is a monthly short course which is a day of intensive talk and interaction. There are weekly courses on demand.

COL. RUSSELL: I would like to point out that as Dave has said many times, ARPA owns this machine, and if you are interested in developing some knowledge about the machine on site, or if you want to interact with the University of Illinois group, we will be more than willing to help you. You can give Dave a call, and I am sure he can make some arrangement for you to visit there, talk to them, see what the problems are, and see how you could best utilize the machine. Is that true, Dave?

MR. MC INTYRE: That is quite true, yes.

COL. RUSSELL: Are there any other comments or questions?

Well, thank you very much.

(Thereupon at 3:10 p.m., the meeting was concluded.)

A SYNTHESIS OF THE PROBLEMS IN SEISMIC COUPLING

William R. Judd
Purdue University

Introduction

These two conferences (June 8-9, 1970: reported in ARPA-TIO-71-13-1, and August 18-19, 1970: reported in ARPA-TIO-71-13-2) established communications between the diverse disciplines required to predict the shock effects from nuclear explosions out to teleseismic distances. These disciplines involve the use of rock mechanics, geology, nuclear physics, computer hardware and codes, seismology, and field instrumentation. Results from the conferences included (a) improvement in the communication links between the engineers and scientists engaged in research relevant to the seismic coupling problems, and (b) identification of open circuits at some points along the communication lines. This paper focuses attention on those open circuits.

In the prototype experiment a nuclear device is embedded in a hole (cavity)* at some specified depth beneath the ground surface. The device is exploded (triggered). The energy produced is partitioned into electromagnetic and radioactive radiation, thermal and mechanical (kinetic) energies. The radiation and thermal energies attenuate rapidly; therefore, their possible appearance at teleseismic distances is ignored. However, the kinetic energy stimulates intense motion of the earth media surrounding the explosion; the resulting body (m_b) and surface (M_s) waves can be identified and measured at distances ranging upwards of thousands of kilometers from the explosion (seismic) source.

This simplified perspective is presented to show why several different scientific disciplines are required to interpret the effects at the measurement point. First, there must be an accurate evaluation of the partition of nuclear energy during and subsequent to the explosion; this quantifies the amount of kinetic energy available to stimulate ground motion. Next, an understanding of how different characteristics of the earth media can affect the propagation of this kinetic energy is required. It is necessary to install instruments that can measure the resulting motions close in to the seismic source. These characteristics and measurements then can be introduced into computer codes designed to describe the orientation and amount of the stresses produced by the ground motion from close in out to teleseismic distances. These stresses can be resolved into the ground displacements that can be expected at teleseismic distances. Measurements are also

---

*There appear to be differences in the use of the word "cavity". Dependent upon the individual user, the word may refer to the hole produced immediately after the explosion, to the hole that develops after the ground in the explosion area reaches stability, or merely to the shape and size of the hole in which the nuclear device is placed.

245

made at teleseismic distances. These are compared with predicted measurements to establish the criteria required to reveal the location and the yield of seismic sources that are inaccessible for U.S. measurements (U.S.S.R. and Communist China).

What Do We Know?

A prominent scientist once said, when discussing the effects of shock waves on hardened installations, that a conference discussing what we know about such effects should be completed within a few hours; however, a conference that discusses what we do not know, would require many days. This philosophy guided the preparation of this report. Part of the conference time was a discussion of what we now can do to predict effects from nuclear devices, particularly at teleseismic distances. The objective was to explain how such effects can be extrapolated to define the yield of explosions that occur in inaccessible areas and also to discriminate between explosions and earthquakes. Our current capabilities in the latter cases had to be qualified by numerous questions relating to the gaps in our prediction ability. This paper summarizes these questions, describes the weak links in the communication lines between the different disciplines involved in the prediction problem, and directs attention to the research required to close the communication gaps.

Role of Geology and Rock Mechanics

If frequent reiteration of a communication problem is any key to its importance, the most significant problem is the lack of numerical methods that will describe the effects of geologic defects, anomalies, discontinuities, etc upon the seismic signal. Time and again the following questions were raised:

"What effect do fractures have upon the energy dispersal and the wave shapes?"

"How can a computer code consider movements along joints?"

"What effect will prestress (also termed 'residual', 'ambient' or 'tectonic' stress) have upon the wave propagation?"

"Can a dispersive model be constructed for jointed and cracked hard rocks?"

"What is the effect of anisotropy in rock properties?"

Ancillary questions were related to the inherent integral properties of a rock element. For example, identification is required of those parameters that can significantly affect either laboratory or in situ tests. Attention has been directed at the changes in wave characteristics produced at various levels of compaction of the rock but there has been little attention to how tensile stresses might

246

affect such characteristics, and, because most waves have a rarefaction phase, it is possible that the behavior of rock in a tensile mode would be of significance.

In Situ Vs Laboratory Properties of Rock

One question that perhaps was most frequently asked was whether the in situ properties of the earth media can be accurately portrayed by laboratory testing. The answers to this question disclosed a divergence of opinion: one group believed it feasible to impose special boundary conditions on the laboratory test specimens to the degree necessary to simulate the prototype performance reliably. However, some conferees felt that reliable answers could be obtained only by in situ tests. A major foundation for these diverse opinions was that because of natural fractures, the in situ media is not a continuum, whereas most laboratory techniques and concommitant analyses are based upon the assumption that the test specimen is a continuum.

Laboratories have used artificially fractured material in an attempt to simulate the effect of joints or fractures. These tests have developed coefficients of friction for such fracture interfaces, but there remains the question of whether such coefficients are valid for natural fractures. Resolution of this problem will require large-scale laboratory or in situ tests. A subsidiary problem is to identify the physical factors that can affect the coefficients of friction on such surfaces.

There also is a need to know the pressures or frequencies or amplitudes that will cause fractures to close and perhaps become transparent to shock waves. Or will discontinuities of this type produce wave refraction and reflection? Most rock systems (and intact rock elements) exhibit some degree of anisotropy in their velocity characteristics, strength, and moduli. There is some evidence that the degree of anisotropy decreases with increasing loads, but further study is required to determine the influence of rock fabric and other natural constituents.

As input to the code calculations it is necessary to have the true in situ compressional velocity, density, isothermal compressibility, water content, compactibility, and the loading and unloading hydrostatic data. At present these values generally have to be obtained or extrapolated from laboratory tests, but their comparison to in situ properties has not been quantified. For example, how does the density determined from an intact laboratory specimen compare with the density of the discontinuous rock system through which the shock-wave propagates? To evaluate the degree of accuracy necessary for such comparisons it will be necessary to conduct parametric studies to define the variation permissible in such values when used in code calculations. A related information gap is the current lack of data on the aforenoted rock properties at pressures up to about 2 kb. There appears to be adequate laboratory data above that pressure level.

A recent step has been taken towards correlation of laboratory and field properties. These studies have found a definite size effect on the Young's modulus of elasticity: the modulus (and the strength) of rock appears to decrease with increasing size of the test specimen. These conclusions are derived from laboratory and in situ tests upon comparable rock elements.

Regardless of the feasibility of achieving a laboratory-in situ test comparability, it was suggested that there would be considerable use for dimensionless rock-property combinations. The latter might provide a more rational method to identify combinations of shot-point rock properties. Also, if such dimensionless values could be established, then instead of using rock names (such as granite, tuff, and alluvium) a dimensionless rock description could be inserted in the magnitude vs yield vs rock-property type of plot. Such dimensionless numbers are difficult to establish because of the wide scatter in the velocity and displacement data that appears to be caused by local cracks, joints, faults, folds, and inhomogeneities. The present analytical approach is to assume a mean value that hopefully will give proper weight to these scatter-inducing properties. The very strong influence of the inherent properties of a rock element has been indicated by the field measurements of such quantities as particle velocity where, at a specific range, such measurements often disagree among themselves by factors of two or three.

Pore Pressure, Porosity, and Water

What are the effects of pore pressure and/or porosity? Does the porosity of a laboratory specimen have a definable relationship to the porosity of the in situ rock system (with its open joints, fissures, etc)? Secondly, how much range or variation in porosity can be tolerated in the code calculation without significant effects on the output? A subsidiary effect of porosity is that an increase in pores may permit an increased water saturation of the material and also a possible increase in pore pressure when the media is subjected to load. The latter occurrence could be of considerable significance in calculations that include media strength because an increase in pore pressure generally means a decrease in effective strength--depending upon whether the pore pressure is sufficient to disrupt molecular bonds between crystals or between grains and the matrix. Another point to be explored in this regard is that in rock (unlike soil) there may be no continuity or connections between pores; therefore, do we have an adequate understanding of the porosity vs water-saturation effects when such rock is subjected to a dynamic load?

The effects of water, including the pore-pressure problem, require considerably more study. There appears to have been insufficient dynamic testing of both intact and cracked material in both the wet and the dry state. Such research is important because it has been established that the change in mass density caused by presence of a water table has an effect upon the wave propagation. A further question

stems from the present assumption that once the depth to the water table is established, all media below that depth must be saturated. Observations in deep tunnels, however, have disclosed tunnel walls that are relatively dry (or, at the most containing only a few percent moisture) even when there are perched water tables above the tunnel elevation. Thus, it is possible that a perched water table might introduce a spurious layering effect in the seismic signatures. There are other possible effects from the presence of water in the media. Relatively close in to the explosion the water may be converted to steam that has an as yet undefined effect on the stress distribution and wave propagation. Also, the effect of water on coefficients of sliding friction between rock elements has not been entirely clarified.

Viscosity

Another factor that appears to have been given too little attention in laboratory and field tests is the influence of the rock viscosity. Theoretically, viscosity should have a strong influence on the high-frequency waves; this has been learned during studies of the transmission of m_b waves in the earth's crust. The effective Q for transmission of m_b waves is on the order of 1000 in the crust but decreases to an order of 100 in the upper mantle. Related factors that may have to be considered in evaluating wave propagation through the crust and upper mantle are the possible movement of interstitial atoms in the lattice, and diffusion of dislocations, partial melt, and pore water.

Failure Criteria

Perhaps the most significant gap in our knowledge of the fundamental properties and behavior of rock is the lack of a reproducible failure criterion. We require a criterion that can provide a mathematical description of the state of the media when failure occurs, including the stress distribution that develops at the failure point. The comparatively recent development of the "stiff" testing machine has made it possible to obtain complete stress-strain curves for many rock materials. For very brittle rock, however, the failure is too rapid to permit delineation of the entire failure path. Therefore, there is a need for a complete stress-strain curve for all rock materials that might house a seismic source.

Reduced Displacement Potential (RDP)

The seismologist measuring effects at teleseismic distances has found that the properties of the earth media definitely influence the reduced displacement potential, but quantification of these effects has not been too successful. The lack of success is attributed to the difficulties in developing a numerical description of geologic defects such as faults, fractures, joints, structure, and stratification.

249

Formulation of a theoretical method that will accurately translate a shock wave from an inaccessible seismic source to a measuring point thousands of kilometers distant presently encounters two major gaps in the transmission sequence: (1) the inability to translate the influence of geologic anomalies into numbers that can be used in code calculations, and (2) the lack of detailed knowledge of the rock properties at the source and between the source and the measurement point. Present opinion is that if we have a geologic description of the earth media at the source we can extrapolate the value of the yield to within 20 to 30 percent of its real value. Also we probably can get within a factor of two of the actual reduced displacement potential if we are provided the density and the seismic velocity of the source material. Our prediction accuracies could be improved if we could establish that the source material had geologic and physicomechanical properties that closely resembled some of the materials intensively studied in field and laboratory tests (such as granite, tuff, and alluvium). However, it was stated that the present dynamic codes might produce a yield prediction that could be in error by a factor of three up to an order of magnitude for such material as tuff! Also, we will require better correlation between the conduct and analyses of nuclear tests and the pre-explosion laboratory and field tests. For example, it was suggested that an objective appraisal be made of the comparisons that have been made between code prediction of nuclear test effects and the actual effects.

Instrumentation and Measurements

Many of our current problems stem from technical deficiencies in our instruments and our procedures. We now lack data on stress conditions at the hypocenters of earthquakes. Therefore we cannot accurately define the resulting seismic-source configuration and establish specific differences between it and a nuclear source. We are severely limited in the depth to which we can make in situ stress measurements. There has been limited success in stress measurements at depths of as much as 4000 ft; however, hypocentral depths are beyond our instrument (and possibly even our drilling) capabilities.

In the laboratory tests, present techniques permit us to measure only the average stress. Thus we must consider the specimen in its entirety; our measurement techniques have not developed to the degree where we can pinpoint the effect of microscopic and, in some cases, macroscopic defects on the stress distribution in the specimen.

One of the most significant gaps in our measurement techniques occurs when we attempt to relate laboratory to in situ measurements. Regardless of whether we are using static or dynamic loading techniques, as discussed previously in this paper, an acceptable correlation between laboratory and field measurements seems to occur as an exception rather than as a rule. Until this gap is closed, we will have to place increasing reliance on field measurements. However this requires us to

develop more reliable and relatively inexpensive methods of making in situ measurements. Also, as was pointed out by one of the conferees, we appear to have no way to make direct use of laboratory-determined material properties to estimate the late-time response of an in situ rock system to an intense shock wave.

Available accelerometers and velocity gages are sufficiently rugged and sensitive to acquire usable information relatively close to the seismic source. However, we do not have a good displacement gage for such close-in effects, particularly one that is capable of measuring displacements on the order of feet in a small-diameter bore hole. At the other end of this spectrum is that because our close-in instruments primarily were designed to measure relatively high motion, they cannot measure strains down to the order of 10^{-5} to 10^{-4}; consequently, in the purely elastic response region such instruments are not effective. We can make reliable measurements at teleseismic distances, but we need a parametric study of instrument capabilities. This may enable the design of instruments having degrees of sensitivity that change with relation to their distance from the seismic source.

Another problem occurs in the establishment of the instrument arrays at teleseismic measurement points. At present, extensive extra-polations of their data are required because only a relatively few instruments are placed at these distances. If we had more stations and azimuth control it could be ascertained whether the geologic structure at the measurement point or the properties of the media at the source control the radiation (of the shock effects) pattern. For example, it would be desirable to have two rings of stations fairly close in to the source and all located within one (geological) structural province where lateral variations in properties were known to be insignificant. Such arrays would permit a study of the radiation patterns as a function of frequency and thus determine whether the theoretical assumptions were correct. The design of such instrumentation, however, necessarily will depend upon a decision as to what parameters should be measured. There are some code specialists who believe that the Rayleigh wave would provide much better information for extrapolation of yield because it samples much more of the structural environment, whereas the P_n wave would not be too good because it considers only a small part of the source region.

One suggested aid to the measurements is to monitor micro-seismic noises in the vicinity of the seismic source prior to the shot. This might provide a clue to the prestressed state of the rock because large stress gradients probably would give a relatively high frequency of noise. At the very least, it would enable a comparison to be made of the ambient stress situations at different shot environments. (Instrumentation for such measurements does exist, and it has been used frequently to monitor potential rock-fall areas in tunnels. Therefore, it merely is a question of adapting this instrumentation for the pur-pose suggested.)

251

The foregoing questions point toward the need for in situ measurement techniques that (a) have a greater reliability than the present ones, (b) can evaluate the changes in properties under dynamic loading, (c) can test several cubic meters of a rock system, and (d) can accomplish the aforenoted measurements without introducing new defects into the rock system. The latter accomplishment would make it possible to test the same rock system under different boundary conditions.

Prediction Code Accuracy

A definitive study of the different codes now used to calculate stress distributions close in to the source indicated that the primary differences between these codes are the manner in which they conserve energy and mass. Some conserve total energy by definition whereas others compute changes in both the kinetic and internal energy analogs and then check each time step to be certain that total energy is conserved to within one part in a very large number (such as 106). Other codes use kinetic and internal energy analogs defined so that the finite-difference equations explicitly conserve total energy.

Teleseismic Prediction

The present codes were designed to study effects close to the source, and they have not been expanded to predict ground-motion effects at teleseismic distances. However, it appears to be within our capabilities to expand these codes so they will produce the latter effects because most, if not all, of the codes now can describe the stress behavior from the source to within the elastic zone. Their expansion to describe effects at teleseismic distances should be relatively simple because the earth media between the present prediction limit and the teleseismic point would be responding as an elastic body.

The first step would be to check the codes for the sensitivity of their calculations. We then could learn what parameters should be measured and just how precise these measurements should be. On the one hand, this will require the seismologists to input the degrees of sensitivity that they require and are able to measure; on the other hand, the rock mechanicist will have to state not only the available sensitivity of laboratory tests but, more importantly, the current capabilities of field instrumentation. For example, is it useful for laboratory measurements to be carried out to one or more decimal places when such precision is not feasible in the in situ measurements? Also codes are structured on the basis that the material being modeled is homogeneous, isotropic, and originally elastic, but, the true media may exhibit none of these properties.

Equations of State

At present we do not know the degree of accuracy required for the Hugoniot data to serve as input for theoretical calculations of the decoupling situation. Our codes also currently presume that we have a complete and accurate equation of state for the earth media subjected to the energy forces. This implies that the equation relates stress, strain, and some of the thermodynamic variables; however, we do not have equations of state for all the types of earth media that might house the seismic source. And, it is not yet clear which rock properties and wave effects are significant in strain-rate dependent behavior. The absence of the latter information makes it impossible to specify the shock-stress levels where purely hydrodynamic rheological effects will occur.

Effects of Heterogeneities and Defects

Existing one-D spherical codes can be used to describe the early stages of ground motion only in a homogeneous media. The introduction of inhomogeneities or defects in the media forces consideration of at least two-D and possibly three-D effects--but such two-D and three-D codes still are in their infancy for such calculations. It was suggested that the calculation difficulty might be partially alleviated if cracks were introduced as an isotropic phenomenon, i.e., they would be assumed to be distributed in such a random manner that there would be no preferential influence on the physicomechanical effects they would produce. However, this introduces the earlier discussed difficulty of defining the wave characteristics at the interface between two cracks. This factor needs resolution, particularly at teleseismic distances where the wave energy is too weak to close the cracks. Thus, in summary, the problem is to determine the degree of wave dispersion close in to the seismic source where the cracks could be closed by the shock energy and the effects at teleseismic distances where dispersed waves would be disrupted further when they encounter fractures that do not close.

Essential to the input of a prediction code are the geologic and rock mechanics data. At present code calculations force-fit preconceived theoretical models for geologic media to the laboratory data, even when there is only a relatively small number of applicable stress states. The requirement is for numerous parametric studies that interface controlled laboratory and field experiments with the code calculations. Such studies would improve the quantitative understanding of the in situ response to dynamic effects.

One empirical finding that has not been predicted successfully by our codes is that the yield vs magnitude curves for different rock types appear to be indistinguishable. For example, unsaturated tuff, granite, and salt all lie approximately on the same curve. Theoretically, the inherent strength of the media elements should exert

an influence on the energy dispersal and thus the inherent physico-mechanical properties of the media should be significant. Resolution of this apparent anomaly would indicate the direction for future research on the rock-mechanics problems associated with nuclear effects. It may be that very close to the source, the rock type is relatively unimportant, but at what critical distance does it become influential, i.e., at what pressure and strain ranges does the rock type become significant? Also it would be of interest to find if defects and inhomogeneities in the rock system exert more influence on wave propagation than do the properties of the intact rock element.

Miscellaneous Considerations

An undecided factor in the calculation of energy dispersion is whether the codes should consider that open fractures may accept large volumes of gas from the explosion. That is, if there are existing fractures or if the explosion opens large fractures, will the latter accept sufficient volumes of gas to attenuate some of the energy relatively close to the source? Our only clue is deductive in that if radiation does not leak to the surface, it is presumed there were no fractures. Part of the answer could be acquired by determining what percentage of the volume of the rock system is occupied by such fractures subsequent to the explosion. Another possible factor in energy attenuation is adiabatic loss. Most of the codes used for ground motion prediction give little or no consideration to such losses because their primary concern is with kinetic energy.

Could codes be made more accurate by decreasing the zoning size, that is, use very fine zoning? It was pointed out that in many cases you would get less accurate answers if this were done, and that for two-D problems, it would not be practical to zone down to a very fine degree. The possibility, however, is that the ILLIAC IV computer may have the capability to handle a very finely zoned problem, particularly those problems derived from two-D or three-D codes.

ILLIAC IV

A brief comment on the ILLIAC IV is appropriate at this point. Most of the conference presentation on this computer related to the hardware although there was a considerable discussion of its operational capabilities. Of special interest to future code calculation is the tremendously increased computation speed as compared with that of existing machines. For example, one of the 64 processing elements in the ILLIAC IV can fetch information from the memory to the operating register in less than one-half the time required by a CDC-6600. Full utilization of the 64 processing elements in the ILLIAC IV will enable it to produce floating point operations at a rate comparable to somewhere between 64 and 128 CDC-6600s.

This new machine should facilitate two-D code work because of the methods it would use to store a matrix and to perform finite-difference calculations. For example, if you want to do one manipulation in the interior of a mesh and a different manipulation on the boundaries, the ILLIAC IV storage capacity and arrangements make it possible to access and parallel all of the values on the top boundary and the bottom boundary because they each are stored in different processing element memories. They then could be copied to the operating registers in parallel and adjustment could be made of the boundary values. Reportedly, these types of calculations may have efficiencies in excess of 80 to 85 percent, i.e., the average number of processing elements turned on during a calculation is approximately 80 percent of 64. Matrix calculation efficiencies generally will be in excess of 50 percent.

On the other hand, accessing information in tables will not be too efficient if the table is so large it cannot be contained in the memory of a single processing element. In particle-motion problems and in nonlinear radiation transport where the particles affect the absorption properties of the media through which they are being transferred, the efficiency may degrade to as low as 25 percent. Another difficulty is that there are no parity checks in the machine at any point. The only way to determine errors is to run a confidence diagnostics program that exercises all of the branches of the logic in the processing element. In other words, you would compute 64 answers simultaneously and determine if any one result differed from all of the others. If so, this presumably would be a logic error.

(NOTE: All of the conferees' statements about the ILLIAC IV were presented prior to actual operation of the machine; presumably, therefore, its precise capabilities and efficiencies are yet to be determined.)

Back to the Codes

A basic and recognized deficiency in code operations is the frequent lack of suitable input data. This deficiency would be alleviated to a considerable extent if there was a comprehensive compendium of all of the test data that is relevant to the calculation of nuclear shock effects. Such a compendium would be particularly valuable if it included time-history details and peak-value tabulations. These would have to be listed in comparison with the more or less standard property data. Such a compendium also would identify significant gaps in the data.

Present codes presume a spherical cavity with a spherical field of motion. Either of these factors can become asymmetric with a resulting degradation in the accuracy of the computation. The amount of such degradation is unknown but it would be desirable to determine the influence of other than spherical cavity shapes and other types of

wave shapes. This could be accomplished by a parametric study designed to evaluate the significance of the resulting differences.

Another useful exercise would be to perform model studies with changing boundary conditions and changing inherent properties. Code calculations then would be performed to see if the results from at least small explosions can be reproduced by codes for various types of materials. The work on just one type of material, tuff, has considered crystal density and porosity but has not introduced water. The latter work is now being initiated, and it is believed that water would introduce a third phase, the first two phases being a porous and a dry material. A related suggestion was to introduce ranges of properties about each main rock type and derive source functions that would correspond to the range of parameters for each particular rock type for a particular yield. This study at least might establish the bounds for the rock types that are studied.

Seismological Input and Output

It would be desirable to modify the codes so they can compute m_b waves and surface phenomena simultaneously with the production of the effects produced by Love waves. One difficulty is that most, if not all of the "large" explosions generate Love waves, but the Love wave does not appear in most lower-yield explosions. Therefore, for code computations using these parameters it would be necessary to define the critical points or boundary lines between yield and the type or types of waves generated vs the distance to the measurement points. And, as stated earlier, the code calculation should be extended to a radial distance sufficient to compute strains as small as 10^{-5}. This would permit a direct comparison between seismological and code calculations.

One point remaining unclarified was whether the present codes can estimate the radial extent of fracturing and crushing out from the source. This definition is required for delineation of the earth-media model that must be used to characterize wave-shape changes and dispersion.

The seismologist would find it useful if the codes could produce the displacement field in potential form within the elastic zone. This implies the definition of ground motions at stress levels of only a few hundred psi, and present codes do not have this capability. The present codes do not contain routines to generate the scalar and vector displacement potentials throughout the region of linear motion. Two-D routines are required and the resulting errors can be on the order of 20 percent or greater. A better feel for two-D problems with a failure mechanism included would permit determination of the true shape of the elastic boundaries around the explosion and in the spall regions. There still would be a need to introduce geological anomalies such as faults, but this might be approached by first doing a calculation that ignores the fault, and then consider the disruptive plane in a manner that

permits an inexpensive parametric approach. For example, the plane could be oriented in various ways to determine the orientation effect on the definition of pressure across the plane. This study could be expanded by evaluating the effect from slip-stick motion and from pre-stress in the media.

The Teleseismic Signature

The seismologist observes a signature on his instruments at the teleseismic recording point--what does it mean? This brings us to the final step in the sequences of wave propagation.

Reduced Displacement Potential

The most important element in an accurate diagnosis of the teleseismic signature appears to be the prediction of the reduced displacement potential of the wave at teleseismic distances. A major control on the nature of RDP is the calculation of the radius at which the earth media starts to react as an elastic body under the influence of the shock. Field measurements and calculations indicate that the RDP is affected seriously by the material properties such as hysteresis and strength. This implies a need to determine late-time displacement in all possible media for all possible source configurations. Although it is known that the RDP is seriously affected by material properties, there is some doubt whether there is sufficient accuracy in the methods now being used to quantify the behavior of these properties. Thus we face the problem of accurately calculating the full range of effects from an explosion close in (where the pressure may be in millions of bars and the temperature in millions of degrees) out to teleseismic distances where the pressures will be a small fraction of a bar and ambient temperatures prevail.

Questions

One diagnostic question is raised by the fact that cavern collapse (at the source) may produce surface waves that appear almost identical to the surface waves produced by the explosion itself; yet the description of these two phenomena in a code calculation would be considerably different. Another question evolves from the situation where the crustal structure at the receiver significantly influences the wave form; therefore it would be desirable to calibrate each source region insofar as the signal level vs yield is concerned.

In general, resolution of the following would assure a better diagnosis of the teleseismic signal and extrapolation back to its source:

(1) How can correlation be achieved between the shot medium and the surface-wave magnitude?

(2) Is it possible to predict which seismic signals in the pass band 0.5 to 2.0 Hz actually propagate out into the elastic zone?

(3) Further attention should be directed to the use of spectral shape as a discriminant although it is recognized that this will not be feasible until there are several azimuths of instrument arrays.

(4) Earthquakes are more efficient in production of Love waves than Rayleigh waves, although the ratio is station dependent and the energy distribution in both time and frequency are different. However, there is not sufficient earthquake data to achieve an accurate diagnosis of the signal by comparing the spectral ratios of Love and Rayleigh waves. Although the exact mechanism of Love wave generation still is unknown, a better understanding might be acquired if theoretical calculations were made near the source. Then, it would be possible and desirable to design a shot that produced propagation effects similar to those from an earthquake.

(5) Can we quantify data distortions that are caused in the short-period data by attenuation, spherical spreading, and layering? The solution of this problem is the key to use of absolute signals as a means of determining the source parameters. There also is a requirement for a model that considers all of the crustal hetero-geneities, including such factors as the variation of velocity and density with depth, the reasons for wave attenuation in different media, and the influence of surface topography, subsurface strati-graphy, and structure. And, although we know that the coupling of energy in hard rock may be an order of magnitude greater than that in soft media, can these distinctions in the source media be identified at teleseismic distances?

(6) The prediction accuracy would be enhanced by efficient operation of two-D codes, including use of a failure mechanism to describe the true shape of the elastic boundaries around the explosion and in the spall region. Surface spall effects clearly are not a linear phenomenon, therefore more precise data is needed on the description of these effects in terms of energy propagating back down into the medium; also, these factors should be expressed as functions of source parameters for a variety of materials.

(7) More accurate predictions would be possible if more precise data were available on the properties of source material that are inaccessible to U.S. investigators. [Author's Note: Such additional data might be extracted from the open Soviet literature on rock mechanics tests within the past decade. This literature rarely indicates the geographic source of the test specimens, but collation of such data may make it possible to group the rock types having similar properties. And, it may be feasible to delete data where the testing evidentially was related to civil, mining, or petroleum engineering projects. Analyses of such collations could provide us with at least a reasonable range of expectable proper-ties in potential source materials.]

(8) There is a requirement for something equivalent to a pressure-time function at a distance where the strains are on the order of 10^{-4} to 10^{-5} and that cover the frequency band of 0.01 to about 2 Hz. Further this pressure-time function should encompass some reasonable volume that encloses the source.

(9) The present codes can predict relative amplitudes of the source, but it is questionable if the codes can provide detailed characteristics of the failure associated with the source. This problem requires knowledge of absolute amplitude and frequency spectra.